

MASTER

Multi-vendor shipment consolidation from shared distribution centers

Schuijbroek, J.M.A.

Award date:
2013

[Link to publication](#)

Disclaimer

This document contains a student thesis (bachelor's or master's), as authored by a student at Eindhoven University of Technology. Student theses are made available in the TU/e repository upon obtaining the required degree. The grade received is not published on the document as presented in the repository. The required complexity or quality of research of student theses may vary by program, and the required minimum study period may vary in duration.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Eindhoven, July 2013

**Multi-vendor shipment consolidation
from shared distribution centers**

by
Jasper Schuijbroek

BSc Industrial Engineering & Management Science
Student identity number 0636544

in partial fulfilment of the requirements for the degree of

**Master of Science
in Operations Management and Logistics**

Supervisors

Prof.dr. Ton de Kok, Eindhoven University of Technology
Prof.dr. Tom van Woensel, Eindhoven University of Technology
Tom Tillemans, H.J. Heinz Company

TUE. School of Industrial Engineering.
Series Master Theses Operations Management and Logistics

Subject headings: transportation, inventory, simulation, management science

Abstract

Vendors co-located in a multi-vendor DC have an opportunity for *multi-vendor shipment consolidation* toward shared customers, but it is unclear which process design ('scenario') should be selected. The most basic scenario is *synchronization* (of order schedules and ship to addresses), which does not require an investment but may yield savings of up to ■% on transportation costs. More integrated scenarios (retailer orchestration and VMI) rely on *coordination* of order quantities with the drop size. The 'push to fill' policy implements such coordination, and simulation shows that it effectively pushes products to fill the vehicles, while improving the service level. However, in The Netherlands the additional savings through coordination do not generate a sufficiently high return on investment due to relatively low transportation costs.

Management summary

Heinz, like many vendors, outsources its outbound warehousing and transportation to a logistics service provider (LSP). Thereby, vendors might become co-located in a ‘multi-vendor DC’. Heinz is co-located with █████ and █████ in a █████-operated DC in █████. Since food vendors mostly sell to the same customers (retailers) and ship to the same retailer DCs (RDCs), this co-location provides an opportunity for ‘multi-vendor shipment consolidation’. However, it is unclear which process design (‘scenario’) is the best choice for different supply chain situations: the *multi-vendor shipment consolidation scenario selection problem*.

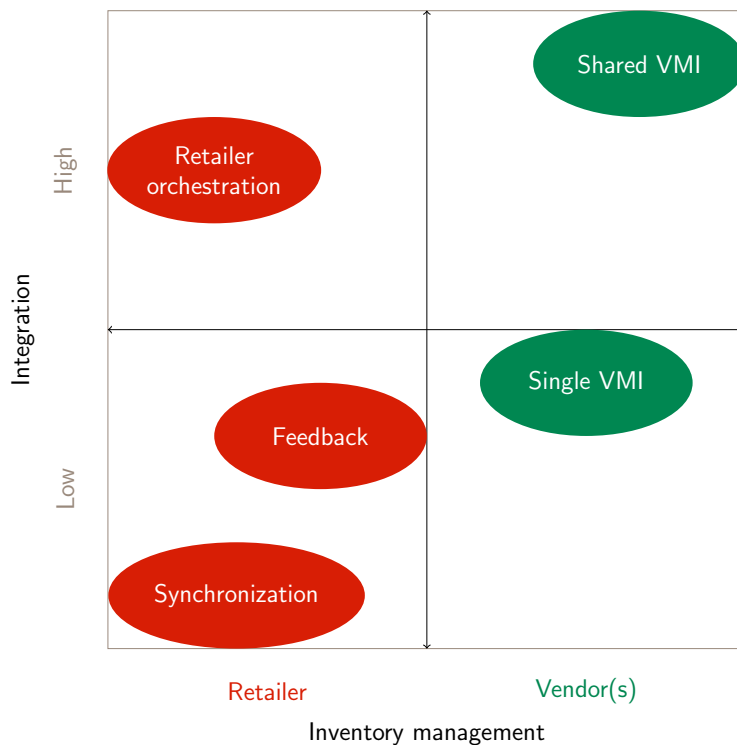


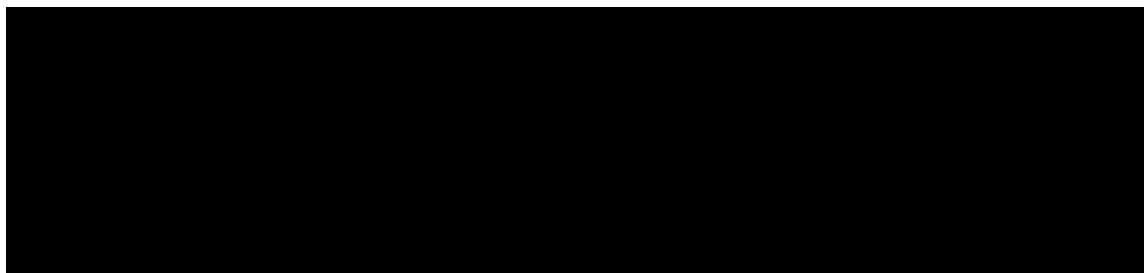
Figure Scenario selection matrix

We identify and introduce five multi-vendor shipment consolidation scenarios and develop a scenario selection framework to support decision makers (see above). Inventory management and the level of integration drive the scenario choice: the more integrated scenarios require higher investments, but might generate a higher return through *coordination* of order quantities with the drop size. We develop the ‘push to fill’ policy which fills vehicles given the ‘pull’ orders resulting from a simple (R, s, nQ) policy. The ‘push

products' are selected based on the lowest run-out time of the reorder level overshoot. The drop size is thereby optimized while improving the service level.

To quantify the cost savings for each scenario, we simulate the push to fill policy in a case study with ██████████. Approximately █████ pallets are shipped from █████ to █████ each year, and the average drop size between November 2012 and May 2013 was only █████ pallet equivalents. █████ shared the historical demand to the retailer DC, hence we simulated the scenarios using realistic parameter settings and extrapolated the results to annual savings¹:

Table Business case



Based on the business case, our recommendation for the ██████████ case is to implement synchronization (of the order schedules and ship to addresses) as soon as possible. This will annually save the vendors €████ (████%) on transportation costs and the retailer €████ (████%) on handling costs. The more integrated scenarios do not yield a sufficiently high return on investment, with shared VMI not even covering its running costs². Once synchronization is implemented, █████ could attempt to prevent the most costly drops (slight retailer orchestration).

Our recommendations for the ██████████ network (vendors and LSP) are:

Improve tariff structure In this case study, we calculated the transportation costs based on a fictional tariff structure. However, in practice the vendors pay different tariffs and are invoiced for their *individual* drop size. Such a tariff structure does not incentivize multi-vendor shipment consolidation.

Experiment with single VMI Multi-vendor shipment consolidation using single VMI is promising, but not as a stand-alone initiative. However, Heinz is discussing VMI with another retailer. This is an excellent opportunity to experiment with pushing products to fill vehicles given the other vendors' drop sizes.

Roll out to all retailers We recommend that the vendors roll out synchronization to all retailers. This could annually save up to €████ on transportation costs and up to €████ on handling costs, without any investment. Moreover, the network should approach large retailers with a business case for retailer orchestration, showing that they can additionally reduce the number of drops by up to █████%.

¹Results should be interpreted in light of the underlying assumptions summarized in Appendix F.

²When we assume that shared VMI can lead to 5% savings for every retailer, which is highly speculative, then the (retailer-specific) running costs of €████ would only be covered for retailers with an annual (synchronized) spend of more than €████.

Preface

Eindhoven, July 2013

This thesis concludes my graduation project at H.J. Heinz Company in Zeist. Heinz is a wonderful company and I would like to thank everyone I've met during the turbulent time the company is going through; especially my direct colleagues Herman Dieleman, Govert van de Heijkant, Karin van der Plas; my co-interns Edwin Hennekes and Lotte Peeters; and my company supervisor Tom Tillemans. Tom, you have an inspirational drive to improve supply chains through network development, and I would like to thank you for the opportunity to work with you and the logistics team.

Next, I would like to thank my university supervisors Ton de Kok and Tom van Woensel. Ton, we've worked together for many years and I'm thankful for learning about your vision on supply chains and the people managing them. Tom, I appreciated your pragmatism and useful feedback.

This thesis also concludes my student life, during which two organizations in particular have been welcome distractions (in the most positive sense!): Industria and the Race of the Classics. I feel happy to have been a part of these great organizations and especially to have shared many memorable moments with the people behind them.

As part of my master's program, I visited the Tepper School of Business at Carnegie Mellon University for one semester. I would like to thank Willem-Jan van Hoeve and Robert Hampshire for inviting and supervising me. I will always remember Pittsburgh and the great people at Tepper. Hope to see you soon!

I would like to thank my parents Jos and Rian, who have always supported me. And finally, having shared most of my student life with Kelly, I look forward to sharing our future.

Jasper Schuijbroek

Contents

Abstract	i
Management summary	ii
1 Introduction	1
1.1 Methodology	2
1.2 Research questions	3
1.3 Company introduction	5
2 Literature review	6
3 Scenarios	9
3.1 Synchronization scenario	11
3.2 Feedback scenario	13
3.3 Retailer orchestration scenario	15
3.4 Single VMI scenario	17
3.5 Shared VMI scenario	19
4 Coordination	20
5 Simulation	24
5.1 Simulating different scenarios	24
5.2 Implementing the push to fill policy	24
5.3 Evaluating results	25
6 Scenario selection	28
7 Case study	30
7.1 Introduction	30
7.2 Problem	32
7.3 Diagnosis	32
7.4 Design	36
7.5 Recommendations	43
8 Reflection	45
8.1 Increasing the shipping frequency	45
8.2 Tariff structure	45
8.3 Effectiveness of coordination	46
9 Conclusions	47

A	Definitions	50
A.1	Sets	50
A.2	Parameters	50
A.3	Functions	51
A.4	Stochastic variables	51
A.5	Decision variables	51
B	Scenario IT landscapes	52
C	Company introductions	56
D	Gain sharing	57
E	Data	58
E.1	LSP data	58
E.2	Retailer data	59
F	Assumptions	61
G	Validation	63
G.1	Model qualification	63
G.2	Model verification	64
G.3	Model validation	64
	Bibliography	68

List of tables

5.1	Push to fill policy for different scenarios	24
7.1	Volume shipped from VDC by vendor	31
7.2	Volume and drop size by address	33
7.3	Volume and drop size by weekday	34
7.4	Demand by weekday	34
7.5	Simulated schedule for both frequencies	37
7.6	Simulation results for drops	38
7.7	Simulation results for inventory	39
7.8	Business case	41
G.1	Verification of simulation in Microsoft Excel ($\mathcal{P}^+ = \emptyset$)	64
G.2	Verification of ‘push to fill’ in Microsoft Excel ($\mathcal{P}^+ = \emptyset$)	64
G.3	Validation results for drops	66
G.4	Validation results for drops (average and standard deviation)	67
G.5	Face validity of simulation results	67

List of figures

1.1	Retail supply chain with multi-vendor DC	1
1.2	Reflective cycle [Van Aken et al., 2007]	2
1.3	Issue tree for multi-vendor shipment consolidation	4
3.1	Synchronization scenario process	10
3.2	Feedback scenario process	12
3.3	Retailer orchestration scenario process	14
3.4	Single VMI scenario process	16
3.5	Shared VMI scenario process	18
6.1	Scenario selection matrix	28
7.1	Supply chain situation	30
7.2	(Fictional) tariff structure	31
7.3	Volume shipped by vendor and period	32
7.4	Histogram of drop size	33
7.5	Costs for distribution of one FTL between two drops	35
7.6	Diagnosis of docking time	36
7.7	Simulation results for drops	39
7.8	Simulation results for inventory	40
7.9	Histogram of drop size for all retailers	42
7.10	Map of ship to addresses of all retailers	43
B.1	Synchronization scenario IT landscape	52
B.2	Retailer orchestration scenario IT landscape	53
B.3	Single VMI scenario IT landscape	54
B.4	Shared VMI scenario IT landscape	55
G.1	Validation of simulation models [Schruben, 1980]	63
G.2	Demand for different instances	65

List of abbreviations

- DC** Distribution center
- EAN** European article number (barcoding standard)
- EOQ** Economic order quantity
- FTC** Freight to customer (i.e. retailer DC)
- FTW** Freight to warehouse (i.e. vendor DC)
- IT** Information technology
- JIT** Just-in-time
- KPI** Key performance indicator
- LSP** Logistics service provider
- RDC** Retailer DC
- RMI** Retailer managed inventory
- ROI** Return on investment
- SSCC** Serial shipping container code
- VDC** Vendor DC
- VMI** Vendor managed inventory
- YOY** Year-on-year

Introduction

This thesis presents the results of our project ‘multi-vendor shipment consolidation from shared distribution centers’ conducted at H.J. Heinz Company (hereafter ‘Heinz’). Heinz, like many vendors, outsources its outbound warehousing and transportation to a logistics service provider (LSP). Thereby, vendors might become co-located in a ‘multi-vendor DC’. Heinz for example is co-located with █████ and █████ in a █████-operated DC in █████.

Since vendors mostly sell to the same customers (retailers) and ship to the same retailer DCs (RDCs), this co-location provides an opportunity for ‘multi-vendor shipment consolidation’: combining ‘drops’ of different vendors to increase supply chain performance (see Figure 1.1). However, multi-vendor shipment consolidation is innovative: the many process designs for consolidation (‘scenarios’) and their impact are not readily available in (academic) literature [Schuijbroek, 2013]. Therefore, this thesis addresses the *multi-vendor shipment consolidation scenario selection problem* (the ‘scenario selection problem’):

Shared distribution centers provide an opportunity for multi-vendor shipment consolidation; but it is unclear which scenario is the best choice for different supply chain situations, both organizationally as in terms of investments and operational performance.

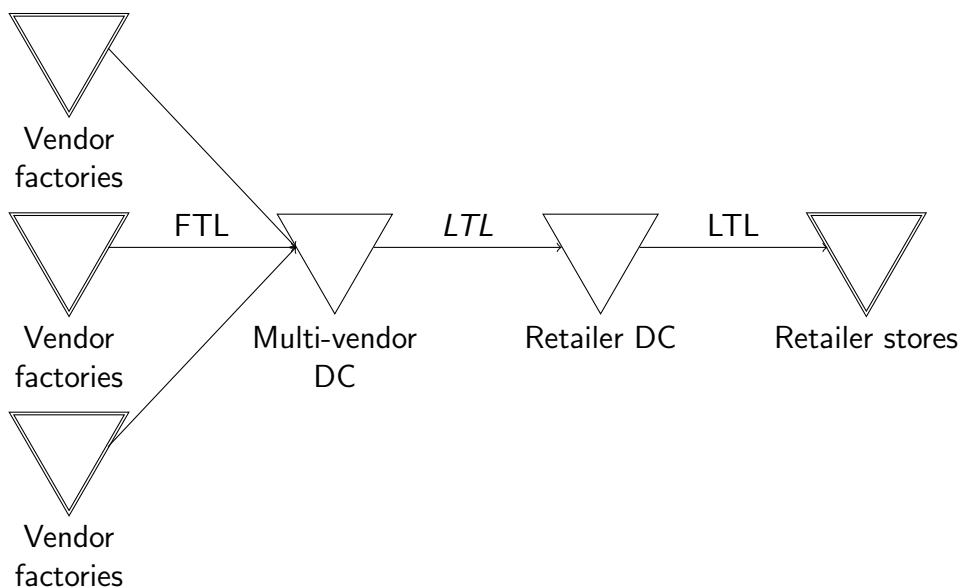


Figure 1.1 Retail supply chain with multi-vendor DC

1.1 Methodology

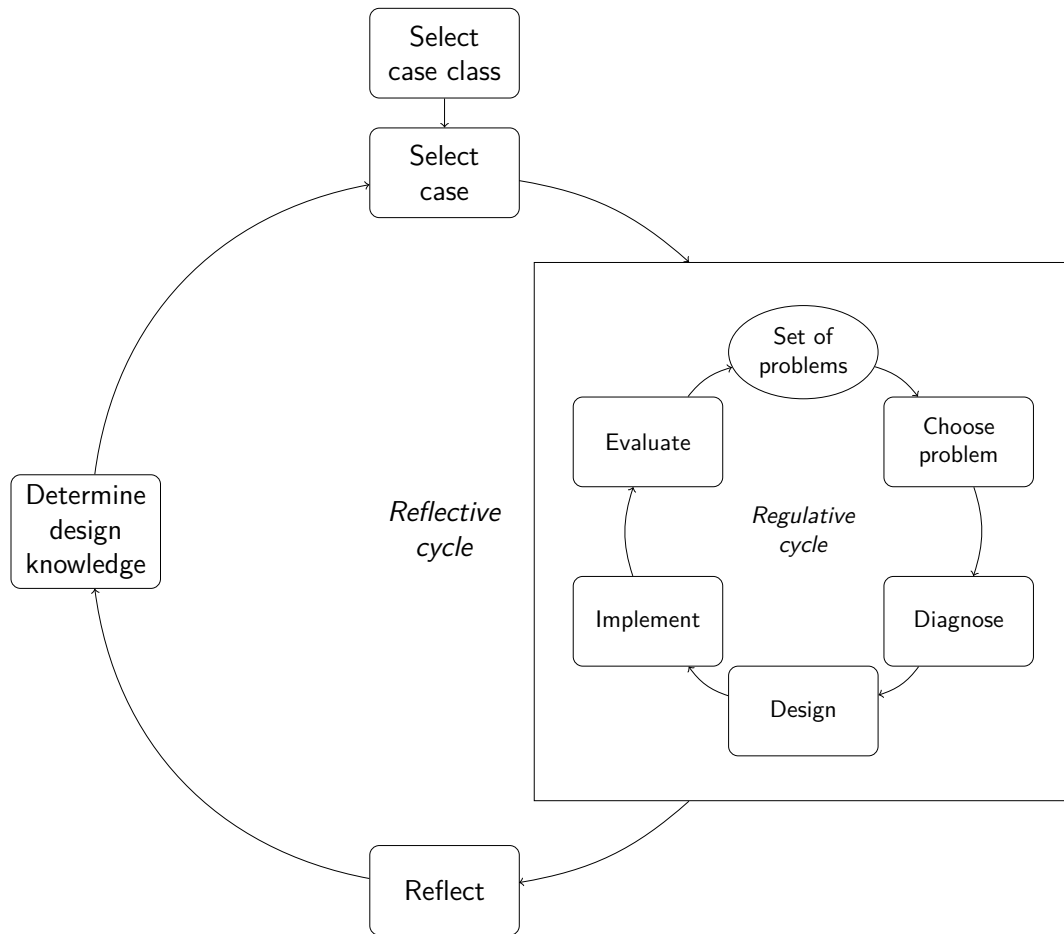


Figure 1.2 Reflective cycle [Van Aken et al., 2007]

We structure the problem solving process (and this thesis) using the *reflective cycle* by Van Aken et al. [2007], see Figure 1.2. Our *case class* is multi-vendor shipment consolidation scenario selection problems.

A multi-vendor distribution center and retailer distribution center define a *case*. Each case is characterized by unique organizational dynamics, an incumbent IT landscape, and complexly interacting operations. Therefore, different cases may require different *scenarios*.

However, we slightly modify the reflective cycle and start by developing generic design knowledge in Chapters 3–6, which we then apply in a case study with █████ (Chapter 7, structured using the regulative cycle). We complete one iteration of the reflective cycle by reflecting in Chapter 8.

The reason for this structural modification is twofold: both the readability of the thesis and Heinz benefit from the separation of *generic* design knowledge and the case study with █████. Heinz serves many retailers from the multi-vendor DC in █████ and operates many (multi-vendor) DCs throughout the world, for which the research questions in the next section are directly relevant, as they are not tailored to the case study.

More practically, our research methods are:

Desk research Both academic literature and best practice documents were consulted to gather design knowledge. These sources are cited throughout the thesis when used. Preparation for the project included a thorough literature review, which is summarized in Chapter 2.

Interviewing The scenario selection impacts all players (vendors, LSP, and retailer), hence (dis)advantages for each player need to be taken into account. Therefore, we performed semi-structured interviews with internal and external decision makers (mentioned when used).

Modeling Since the design phase of the regulative cycle should be supported quantitatively, we develop a conceptual model in Chapter 4 and Chapter 5.

1.2 Research questions

We define five research questions to structure the development of generic design knowledge in the first five chapters (3–6). The answers to these research questions are summarized in Chapter 9.

1. What are the (dis)advantages of multi-vendor shipment consolidation?
2. What are feasible multi-vendor shipment consolidation scenarios?
3. How to coordinate multi-product inventory and shipment decisions?
4. What is the operational performance of the different scenarios?
5. How to select a multi-vendor shipment consolidation scenario?

We hypothesize about the (dis)advantages of multi-vendor shipment consolidation in Figure 1.3, which is based on extensive interviews with experts (see Chapter 3). The two main advantages of multi-vendor shipment consolidation should be a useful increase in the shipping frequency and an increase in the shared drop size [Coppens, 2012]. We answer the first research question by subjecting these hypotheses to quantitative analysis in Chapter 7 (see methodology above).

Second, in Chapter 3 we attempt to identify the different scenarios for multi-vendor shipment consolidation and to provide generic knowledge on their processes, IT landscapes, organizational dynamics and operational performance.

We quickly find out that multi-vendor shipment consolidation beyond ‘synchronization’ requires coordination of order quantities (inventory) and drop sizes (transportation) [cf. design choice 1 in Van der Vlist, 2007, p. 9]. In Chapter 4, we attempt to answer the third research question by developing a new ‘push to fill’ policy which optimizes the drop size given inventory requirements.

Fourth, we develop a simulation model in Chapter 5 to quantify the operational performance of each scenario. Finally, in Chapter 6 we develop a scenario selection framework to answer the fifth research question.

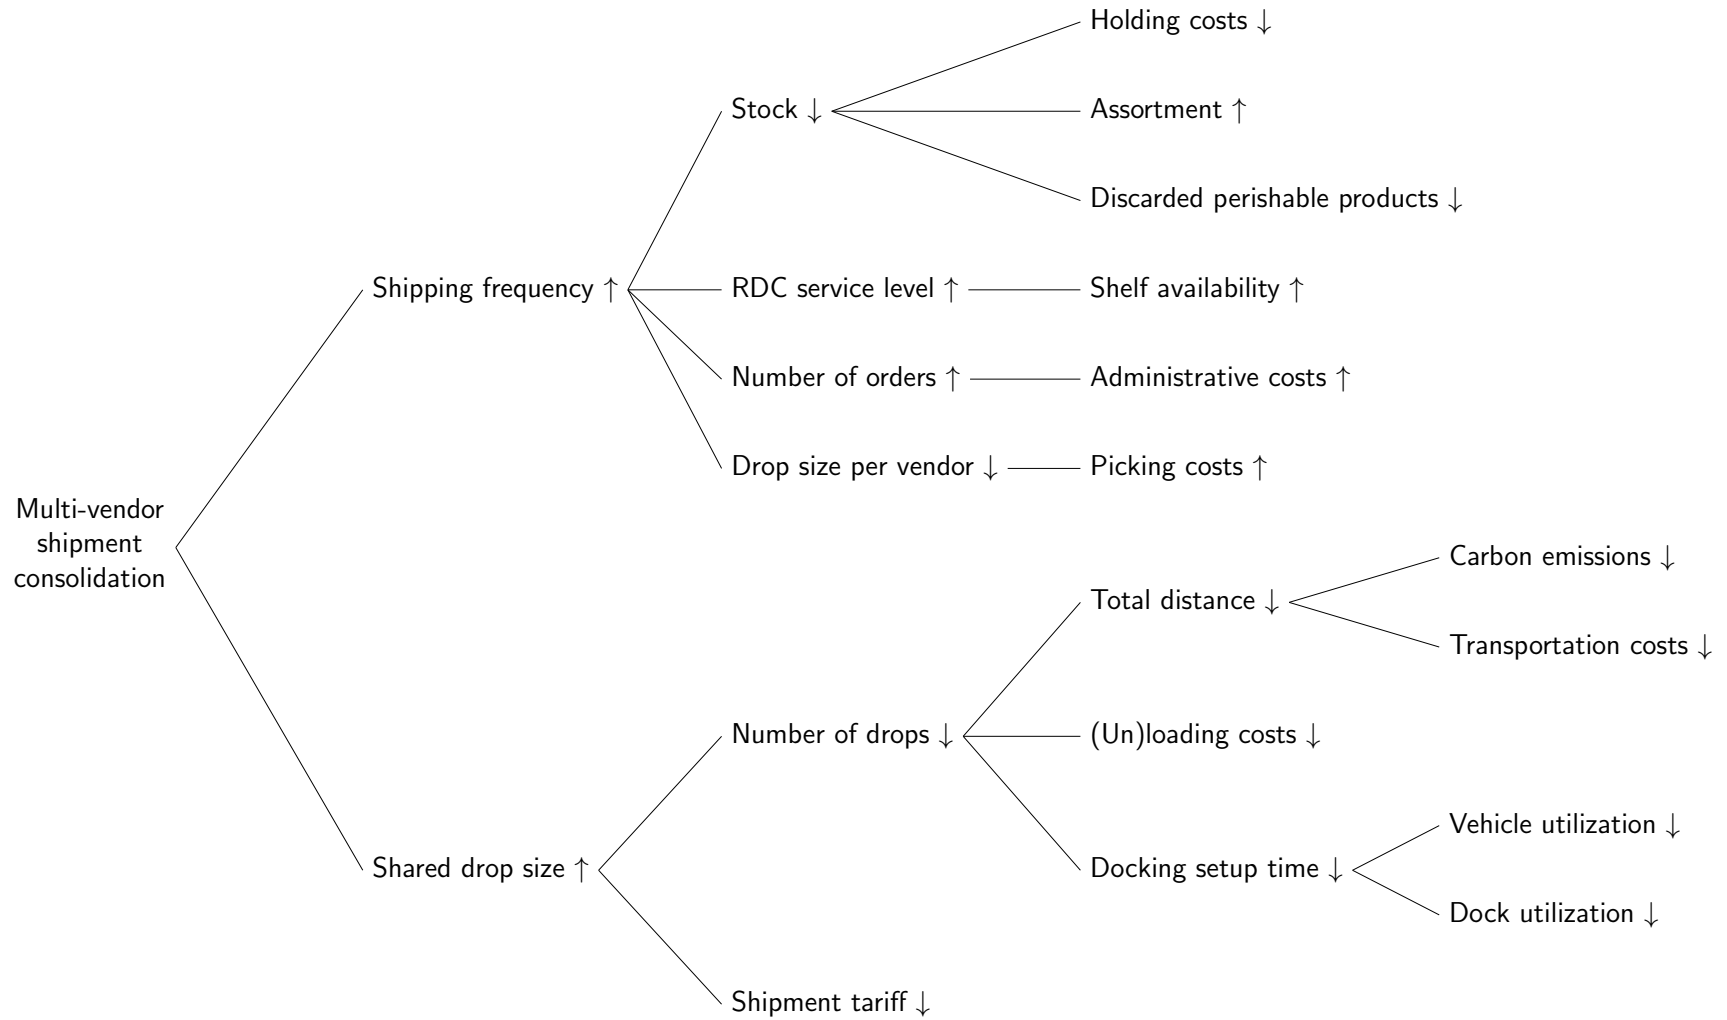


Figure 1.3 Issue tree for multi-vendor shipment consolidation

1.3 Company introduction

Two

Literature review

In preparation for this thesis, relevant academic literature was reviewed in Schuijbroek [2013]. Below is a short summary.

Shipment consolidation

Even for a single vendor, (analysis of) multi-*product* shipment consolidation is complex. Therefore, the literature review first focuses on shipment consolidation: how to create full truckloads (freight packing), how to coordinate inventory with transportation, and how to combine orders (consolidation strategies).

Freight packing

In a single-product environment, a full truckload shipment contains a known number of products. However, food retail is inherently a multi-product business. Freight packing focuses on maximizing capacity utilization *given* a set of orders to be shipped.

The computationally complex (one-dimensional) *bin packing problem* requires a set of heterogenous items to be packed in a minimal number of finite-capacity bins [Johnson, 1974]. Note that due to the complexity, *complete* usage of a bin is uncommon. Since freight packing of trucks is a three-dimensional bin packing problem with additional constraints [ORTEC, 2013], the strong complexity forces the use of simplifying heuristics in practice.

The simplest heuristic is the use of the conversion factor $\alpha_i \in (0, 1]$, which approximates the capacity requirement of a single product $i \in \mathcal{P}$ [Coppens, 2012, p. 28]. It is assumed that if the total capacity requirement is smaller than or equal to the total capacity of x vehicles ($= Cx$), then there exists a feasible assignment of the products to these vehicles. Unfortunately, this approach has not been validated (neither in practice nor in literature), and since Coppens focuses on a simple case ($\alpha_i = 1, \forall i \in \mathcal{P}$), this is mentioned as a direction for future research [p. 49].

Coordination

The interaction between transportation costs and inventory management has been studied before. According to Higginson [1995, p. 5] most approaches are ‘economic shipment quantity (ESQ)’ models, based on the single-product economic order quantity (EOQ):

$$\text{ESQ} = \sqrt{\frac{2\lambda C}{h}}$$

with order arrival rate λ , fixed vehicle dispatch costs C and holding cost rate h . A natural alternative is the economic shipment weight (ESW), for which λ and C are weight-based. Note that in practice, the ESQ is bounded from above by the vehicle capacity.

Swenseth and Godfrey [2002] are the first to address the optimal order quantity of a single product given a tariff structure for transportation. They present an interesting insight into overdeclaring weights and provide a strong example of the suboptimality of an EOQ-based approach. These results should be extended to the multi-product case (see Chapter 9).

Consolidation strategies

However, coordination assumes centralized shipment and inventory decisions (e.g. using vendor managed inventory). In food retail, retailer managed inventory is standard, and vendors are not allowed to ship more than retailers have ordered. This implies that FTL shipments can only take place once retailer orders exceed the vehicle capacity. Therefore, vendors use different *consolidation strategies* (or dispatch policies) to increase the drop size. These policies simply pertain to the *timing* of shipping retailer orders:

Time-based consolidation strategy Vendors postpone dispatch of an order up to a time limit T to accumulate additional orders. Order schedules can be viewed as an implementation of the time-based consolidation strategy.

Quantity-based consolidation strategy Orders are accumulated and a vehicle is not dispatched until a set quantity Q is reached [Cachon, 2001, p. 212].

Time-and-quantity-based consolidation strategy Orders are accumulated until either a time limit T or quantity Q is reached [Mutlu et al., 2010].

Multi-vendor collaboration

Next, Heinz's co-location with other vendors yields opportunities for multi-*vendor* shipment consolidation. We address both the quantitative aspects (benefits and gain sharing) and qualitative aspects (organizational challenges) of multi-vendor collaboration.

Benefits of multi-vendor shipment consolidation

Coppens [2012, p. v] establishes that the two main benefits of multi-vendor shipment consolidation are:

Useful increase in shipping frequency The shipping frequency can increase if a vendor $v \in \mathcal{V}$ enters into a coalition with co-located vendors $\mathcal{V} \setminus \{v\}$.

The useful shipping frequency in the coalition is larger than or equal to each vendor's individual shipping frequency. We note that the increased shipping frequency is only useful for products with a high enough demand $E[D_i]/Q_i$ due to batching.

Higher shared drop size If a vendor is already using consolidation strategies effectively to achieve a close-to-optimal drop size, then the increased useful shipping frequency is the only benefit of multi-vendor shipment consolidation. However, we note that most vendors are shipping at much lower utilization rates.

Gain sharing

Gain sharing is a well-studied sub-discipline of *cooperative game theory*. A *transferable utility (TU) game* is a pair (\mathcal{V}, ν) with value function $\nu : 2^{\mathcal{V}} \rightarrow \mathbb{R}^+$ assigning a value to each coalition $S \subseteq \mathcal{V}$ with $\nu(\emptyset) = 0$ [Slikker, 2012, p. 71]. The *core* consists of the efficient¹ payoff vectors which satisfy stability, $\sum_{v \in S} z_v \geq \nu(S)$ for all $S \subseteq \mathcal{V}$. In a TU game with non-empty core, no (sub-)coalition has an incentive to leave the grand coalition \mathcal{V} .

One provable core element (i.e. stable payoff vector) is the *Shapley value* [Slikker, 2012, p. 87]. Let $\sigma : \mathcal{V} \rightarrow \{1, \dots, |\mathcal{V}|\}$ be a bijection indicating the order of the players of coalition \mathcal{V} . The Shapley value Φ_v is equal to the expected marginal contribution of a player $v \in \mathcal{V}$ under a randomized entrance ordering:

$$\Phi(\mathcal{V}, \nu) = \frac{1}{|\mathcal{V}|!} \sum_{\sigma \in \Sigma(\mathcal{V})} \nu'_\sigma$$

with ν'_σ the marginal vector under ordering σ and $\Sigma(\mathcal{V})$ the set of possible orderings of \mathcal{V} . The Shapley value satisfies efficiency, additivity, the zero-player property (zero-player receives no payoff), and symmetry (identical players receive identical payoffs).

We note that calculating the Shapley value requires knowing the value of exponentially many (sub-)coalitions. Computation of these coalition values is often complex. Therefore, approximations may be used, e.g. continuous approximations for transportation costs [Langevin et al., 1996; Daganzo, 1996].

Rules of thumb for dividing the gains are mentioned as a practical alternative to the formal allocation methods from the previous section, e.g. proportional to the total load shipped [Cruijssen et al., 2007, p. 31]. However, ‘in the long run, some participants will inevitably become frustrated since their true share in the group’s success is undervalued’ [p. 32]. Any fair allocation should be based on the marginal contribution of a participant to the coalition.

Organizational challenges

Confirming the existence of a profitable coalition and a ‘theoretically’ fair allocation of benefits and costs is not sufficient to implement change. Many organizational challenges arise in (horizontal) collaboration.

Simatupang and Sridharan [2002, p. 21] present six interventions for effective (vertical) supply chain collaboration: mutual objectives, integrated policies, appropriate performance measures, decision domain, information sharing, and incentive alignment. For each of these interventions, the authors discuss short-, medium-, and long-term benefits. Short-term incentive alignment strategies are ‘productive behavior-based incentives’ and ‘pay-for-performance’ [p. 21]. Medium-/long-term incentive alignment should be based on ‘equitable compensation’ with ex-ante mechanisms and an ‘open book practice’ [p. 27].

As mentioned by Coppens [2012], horizontal collaboration cannot succeed without vertical collaboration. Mason et al. [2007] analyze this topic in more detail. It is important to draw in competitors (horizontal) as well as suppliers and customers (vertical) [p. 193]. They present three case studies motivating the necessity of vertical collaboration.

¹An efficient payoff vector pays out all value: $\sum_{v \in \mathcal{V}} z_v = \nu(\mathcal{V})$.

Three

Scenarios

In this chapter, we present the different scenarios for multi-vendor shipment consolidation. For each scenario, we shortly discuss the process¹ (improvements), IT landscape, and organizational dynamics. The operational performance of each scenario is highly dependent on the case data, but we provide some generic remarks.

From interviews with experts, and best practices, we created five scenarios for multi-vendor shipment consolidation:

- Synchronization;
- Feedback;
- Retailer orchestration;
- Single VMI;
- Shared VMI.

These semi-structured interviews were conducted with the ‘customer service manager’ and ‘customer service logistics manager’ at Heinz, the ‘business development manager’ at ██████ (an LSP), and the ‘logistics project manager’ at ██████. ██████ had already developed an overview of scenarios, but process designs were out of scope. In addition, best practices from Heinz ██████ (shared VMI scenario) and retailer ██████ (feedback scenario) were incorporated. The scenarios were discussed and agreed upon in a presentation with representatives of all players in the case study (see Chapter 7).

¹Note that all ordering processes are simplified, since for effective inventory management, not only the inventory position but also forecasts, discounts etc. need to be taken into account. For VMI scenarios, such information should be clearly communicated by the retailer.

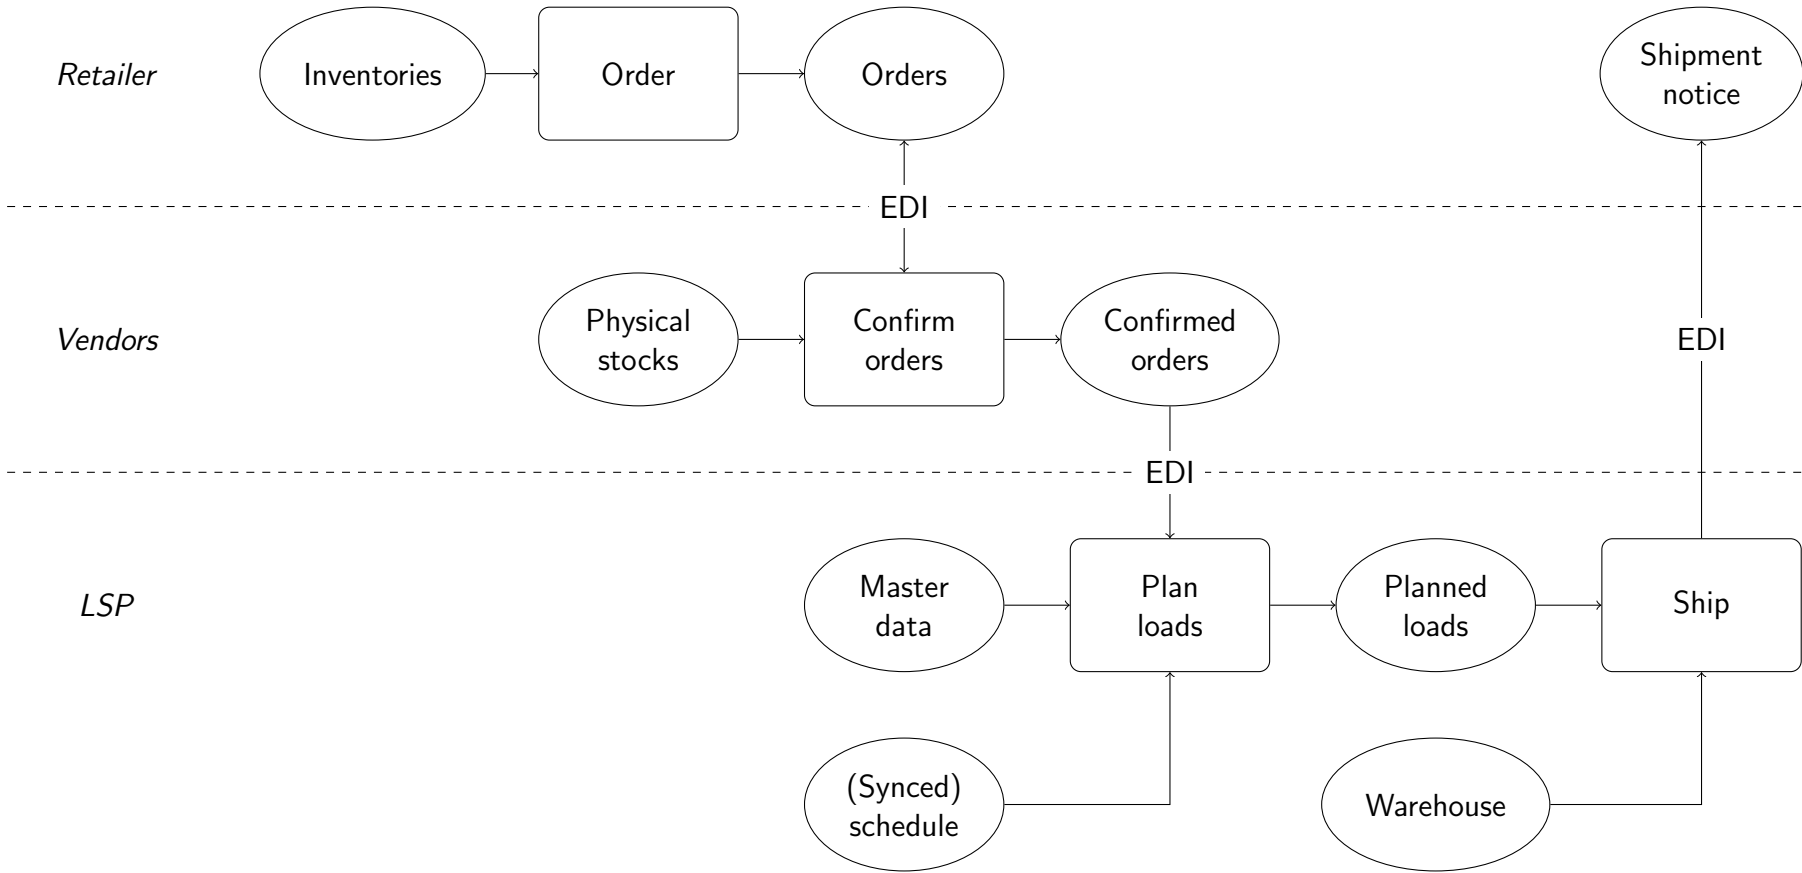


Figure 3.1 Synchronization scenario process

3.1 Synchronization scenario

Synchronization is the most basic scenario and does not require process changes. We show the (original) process in Figure 3.1 for reference. Purely by synchronizing delivery schedules between vendors (i.e. shipping to retailer distribution centers on the same days and times), benefits may be obtained, since the LSP will consolidate shipments to minimize logistics costs.

The process shows that orders act as ‘constraints’: once confirmed, they must be shipped at all costs, which can lead to vehicle underutilization. Further improvement, as implemented in the remaining scenarios, requires coordination between retailer inventory (orders) and transportation, which will be discussed in Chapter 4.

IT landscape

Synchronization can be implemented without modifications to the incumbent IT landscape, which is shown for reference in Figure B.1 on page 52. In general, the incumbent IT landscape is based on schedules and synchronizing these schedules does not require new IT.

Organizational dynamics

Since no investment or further integration is required to implement synchronization, the organizational dynamics will most likely be smooth.

We note that the tariff structure should incentivize multi-vendor shipment consolidation: vendors paying tariffs related to the collective drop size, instead of their individual drop size (which will decrease as shown in Figure 1.3). This is a form of ‘benefit sharing’ by the LSP.

This might lead to (large) vendors complaining about the fairness of a shared tariff structure when volumes are significantly different between vendors. However, this can be solved by differentiating the pallet equivalent tariff between vendors while measuring the shared drop size.

Operational performance

Co-located vendors should always implement synchronization (i.e. synchronize their schedules). Even without coordinating order quantities, this will never decrease performance, since synchronized ordering policies can be parameterized to replicate each vendor’s pre-synchronization shipments.

We note that the operational performance of the synchronization scenario is ‘hit-and-miss’. If the combined orders under a synchronized delivery schedule can consistently be shipped in FTLs, then synchronization is a highly efficient scenario.

However, variation in orders may lead to costly overshoots of vehicle capacity [Coppens, 2012]. In addition, the performance of the synchronization scenario is highly sensitive to changes in demand and in the group of vendors.

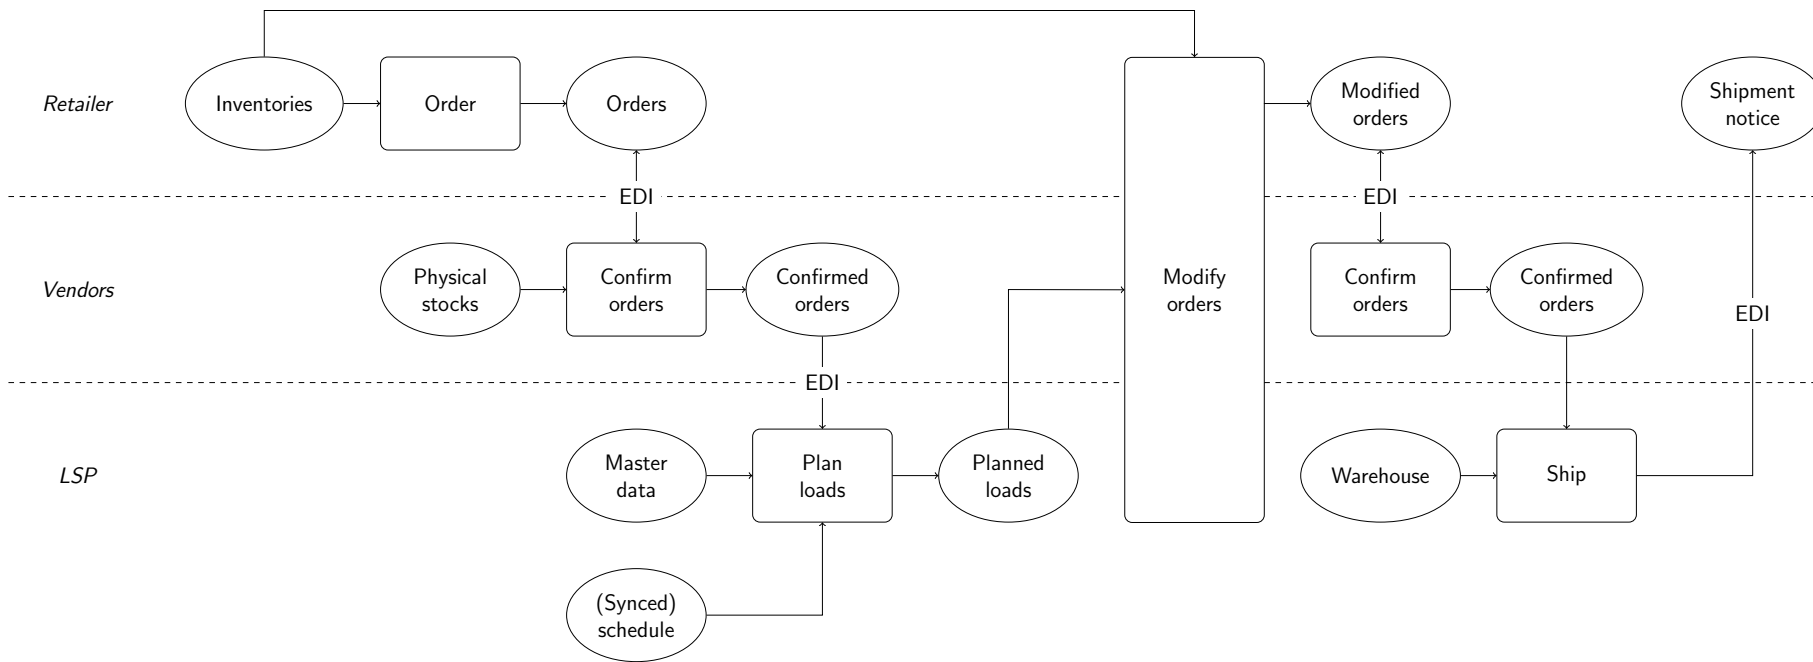


Figure 3.2 Feedback scenario process

3.2 Feedback scenario

The feedback scenario shown in Figure 3.2 extends the synchronization scenario (Figure 3.1) with a feedback loop in which the players may adjust orders based on feedback from the LSP. The vendor is involved for two reasons: incentives to improve the drop size (the vendor pays) and agreements on the drop size with the retailer (e.g. in exchange for a certain schedule). The disadvantage of such a feedback loop is the increased workload for which administrative costs are incurred.

We identify two special cases of the feedback scenario:

Order modification Last-minute drop size modification means that the LSP adds or removes products at the dock, which is highly effective operationally, but requires extreme IT agility. For customer service and data consistency reasons, the retailer has to add or modify orders at the source.

Lead time modification Alternatively, the LSP might choose to ship *ordered* products earlier than intended due to slack capacity.

IT landscape

Depending on the implementation, the feedback from LSP to vendor(s) may be either ‘manual’ (e.g. by telephone) or automated through an IT interface (see Figure 3.2). Currently, Heinz is experimenting with the manual version of this scenario in cooperation with food service customer ██████████.

Organizational dynamics

The feedback loop requires strong relationships between the operational workforce of the LSP, vendor(s), and retailer to ensure a swift execution. The time window for improvement is often only several hours long.

Operational performance

Naturally, the feedback loop will fix the worst shipment errors (i.e. the most costly overshoots). We argue that due to the lack of sophisticated (load building) software to handle the product master data at the vendor and retailer, there will always remain some undershoot of vehicle capacity.

However, the lack of significant investments may make this scenario attractive in cases with low demand variance for which merely synchronizing schedules is insufficient.

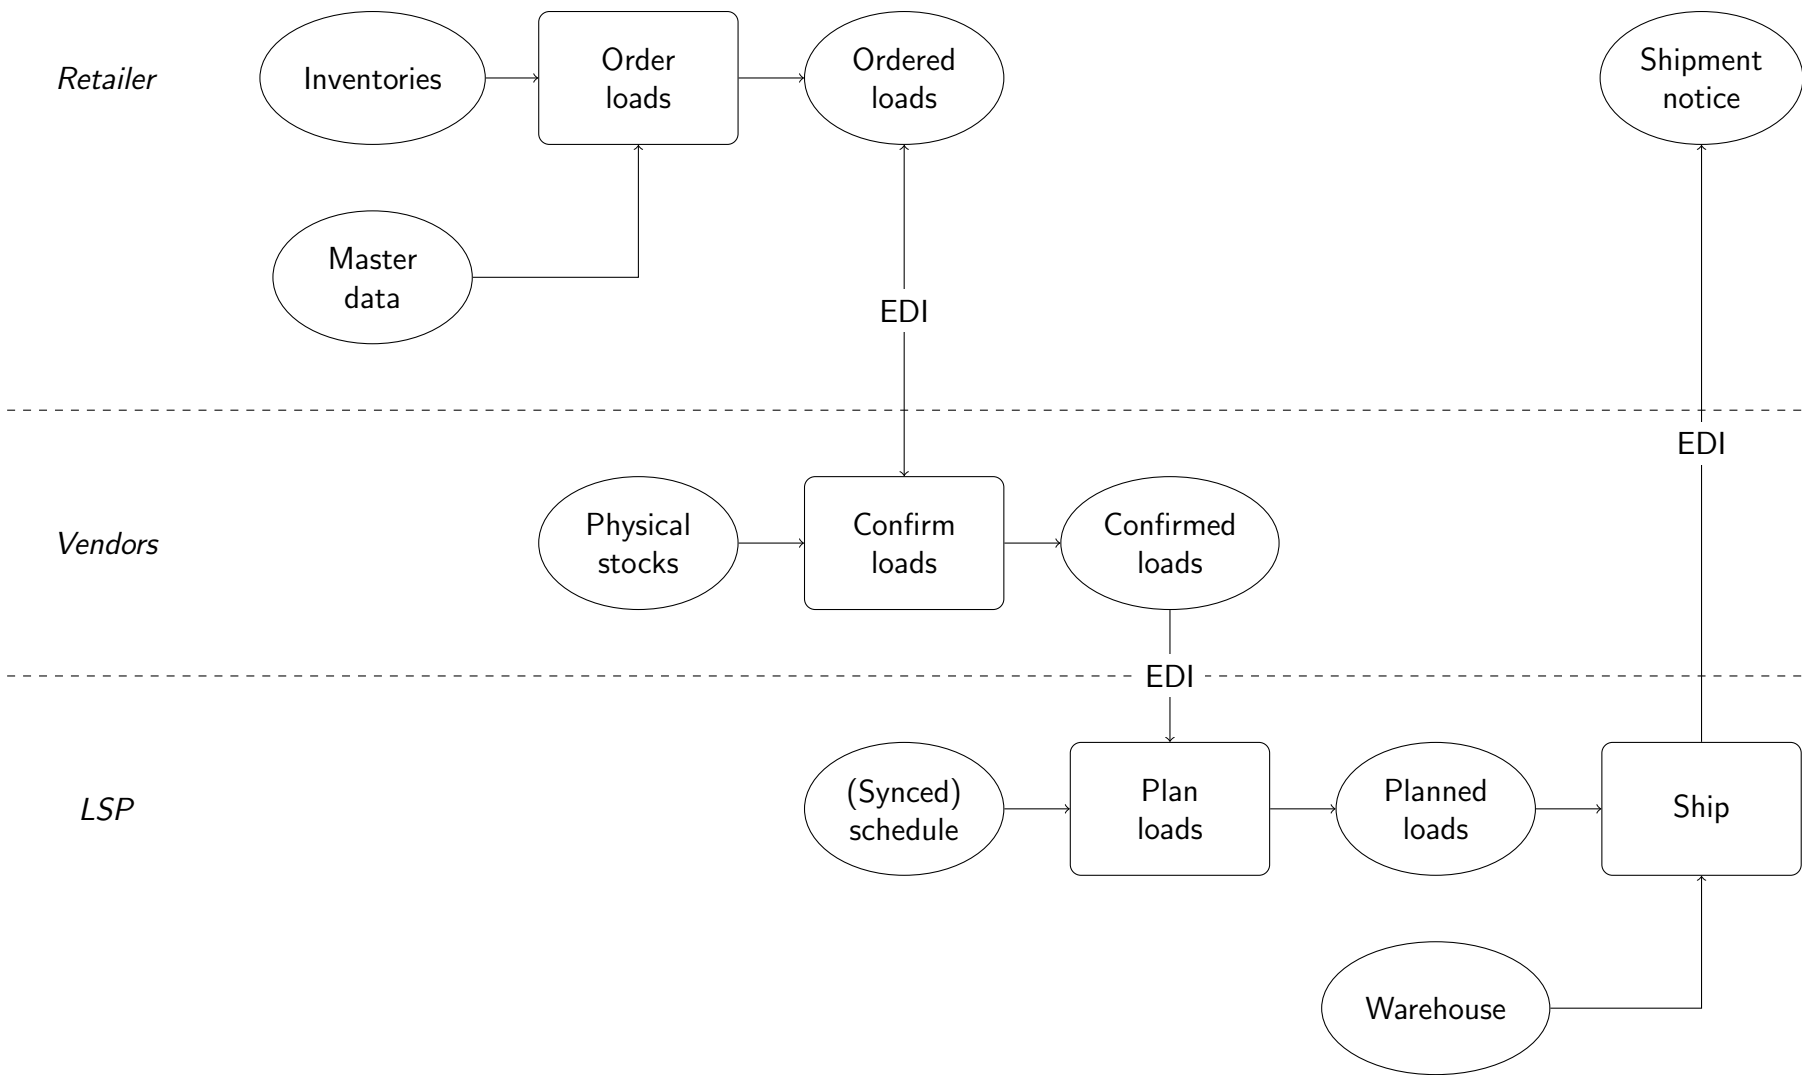


Figure 3.3 Retailer orchestration scenario process

3.3 Retailer orchestration scenario

Figure 3.3 shows the retailer orchestration scenario, which only extends the synchronization scenario by enabling the retailer to build loads with product master data from the vendors.

IT landscape

Load building software is necessary to enable the retailer to order efficient ‘loads’ instead of products. ORTEC provides such a load builder, which takes into account dimensions, weight, interstackability, etc. [ORTEC, 2013]. This software retails for approximately €■■■.

The load builder is shown dotted in the IT landscape in Figure B.2 on page 53: using iterative calls to the load builder, orders can be added by the ordering software until the vehicle utilization is maximized (or optimized in conjunction with stock levels).

Organizational dynamics

Since the retailer is clearly driving improvements in this scenario, we can expect interesting organizational dynamics. We believe the retailer’s investment in IT can only be profitable when the scenario is implemented for several vendor DCs (i.e. cases). Therefore, the retailer is probably the initiator and may demand a significant part of the savings.

Furthermore, there is no requirement for a strong relationship between the vendors. Each vendor should individually make an effort to supply the retailer with detailed product master data, which requires the retailer to be a sufficiently sized customer to each vendor. Since the vendor usually pays for freight to customer (FTC) transportation, the incentive for sharing such data is in place. Note that this scenario would be especially relevant when the retailer would pay for transportation [factory gate pricing, see Le Blanc et al., 2006].

Operational performance

In some sense, this scenario is ‘utopia’, since the coordination problem of inventory & transportation is solved at the origin. However, from the process diagram we learn that out of stock situations at the vendor may lead to underutilization of vehicle capacity. The vendor is unable to share such physical stock information, because the vendor needs flexibility in case of shortage allocation between multiple retailers (‘commonality’ of the physical stock). Depending on the vendors’ DC service level, this could decrease operational performance².

²The drop size does not suffer from out of stock situations at the multi-vendor DC in vendor managed scenarios.

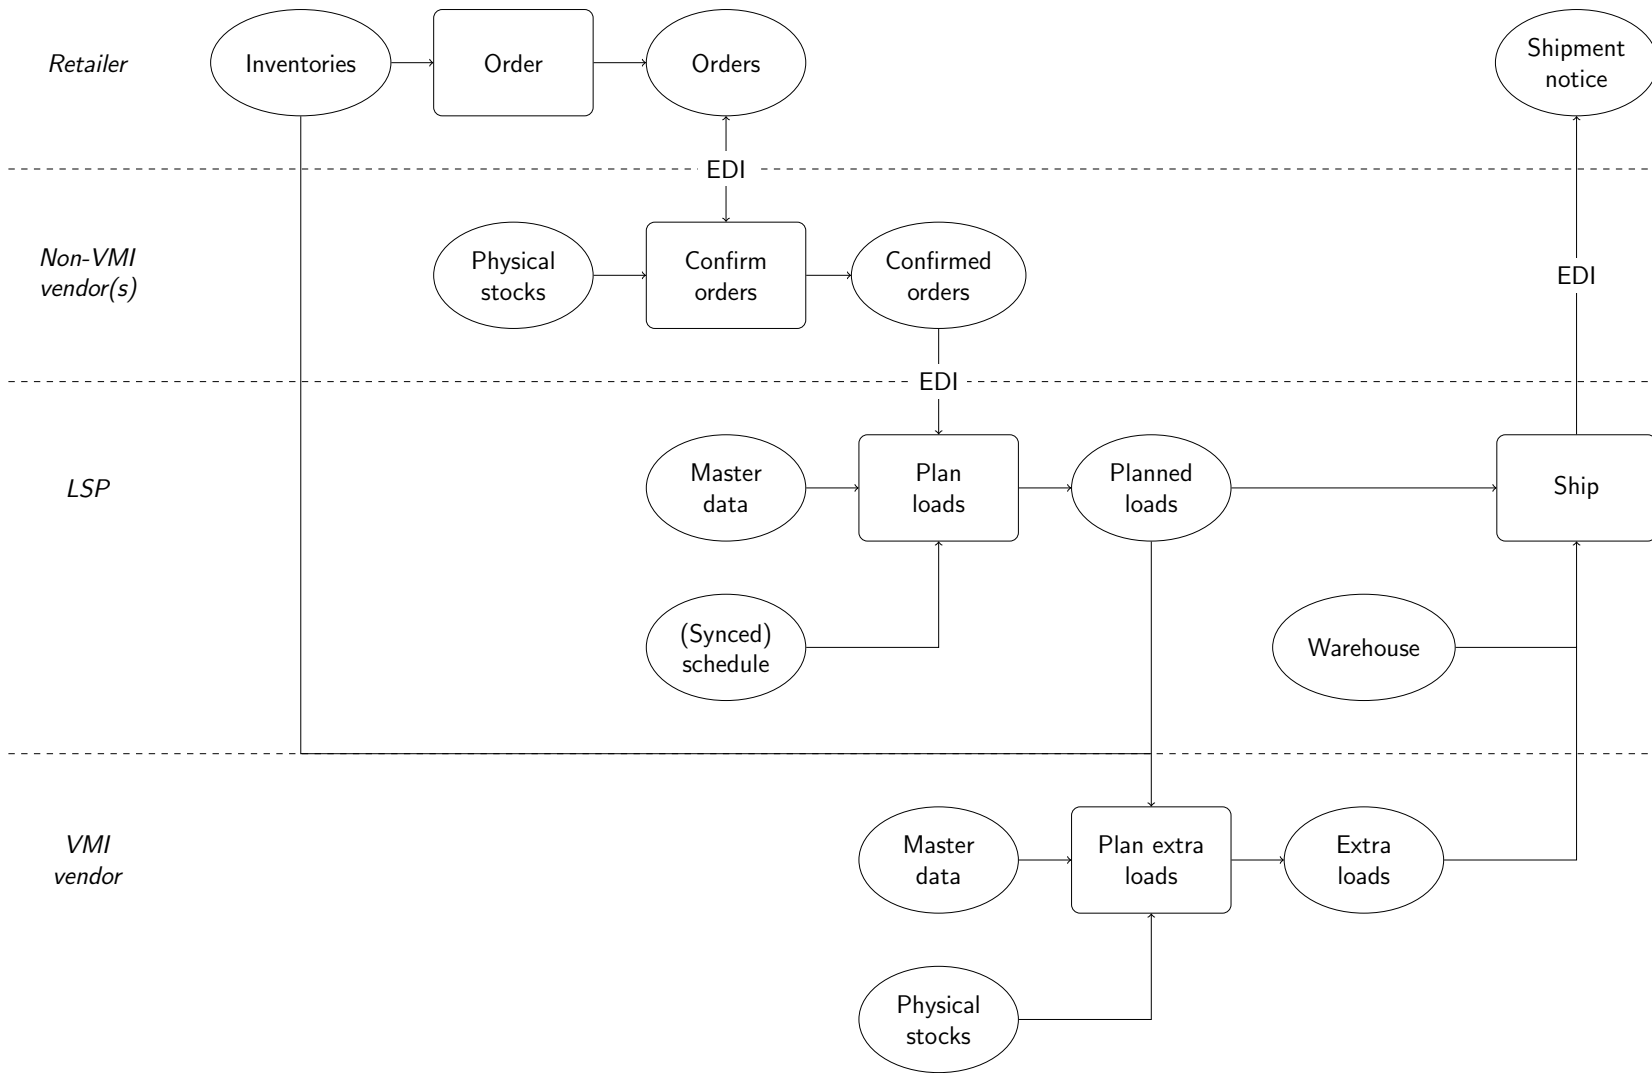


Figure 3.4 Single VMI scenario process

3.4 Single VMI scenario

The single VMI scenario (see Figure 3.4) extends the synchronization scenario by having one of the vendors implement VMI. In the food supply chain, VMI means that a vendor manages the inventory at the retailer (i.e. releases orders), often subject to a minimum and/or maximum inventory level, or to service level requirements. After the LSP plans confirmed retailer orders at the non-VMI vendors, the VMI vendor receives load information and may decide on its shipments given those for other vendors (and fill the vehicles).

IT landscape

The IT landscape for single VMI is shown in Figure B.3 on page 54. Depending on whether the VMI vendor has already implemented VMI, the appropriate interfaces with the retailer should be created. In addition, standard VMI software may not be able to consider shipment efficiency, especially when filling vehicles given shipments for other vendors, in that case a load builder is necessary. This might require some significant investments (similar to retailer orchestration, approximately €■■■■).

Organizational dynamics

The organizational dynamics of the single VMI scenario can be interesting. Since only one vendor is filling the vehicles:

- The non-VMI vendors could scrutinize the VMI vendor for low performance;
- The VMI vendor could perceive unfair gain sharing, since he is ‘doing all the hard work’.

A combination of contractual agreements and strong network-oriented relationships can mitigate such issues.

Operational performance

Selecting the appropriate VMI vendor is crucial to achieve significant improvements. Since filling vehicles may sometimes require a high number of pallets, the VMI vendor should be one of the larger players in terms of volume. Moreover, a strong relationship with the retailer helps: the VMI vendor will sometimes ship more than strictly necessary, which leads to higher stocks in the retailer DC.

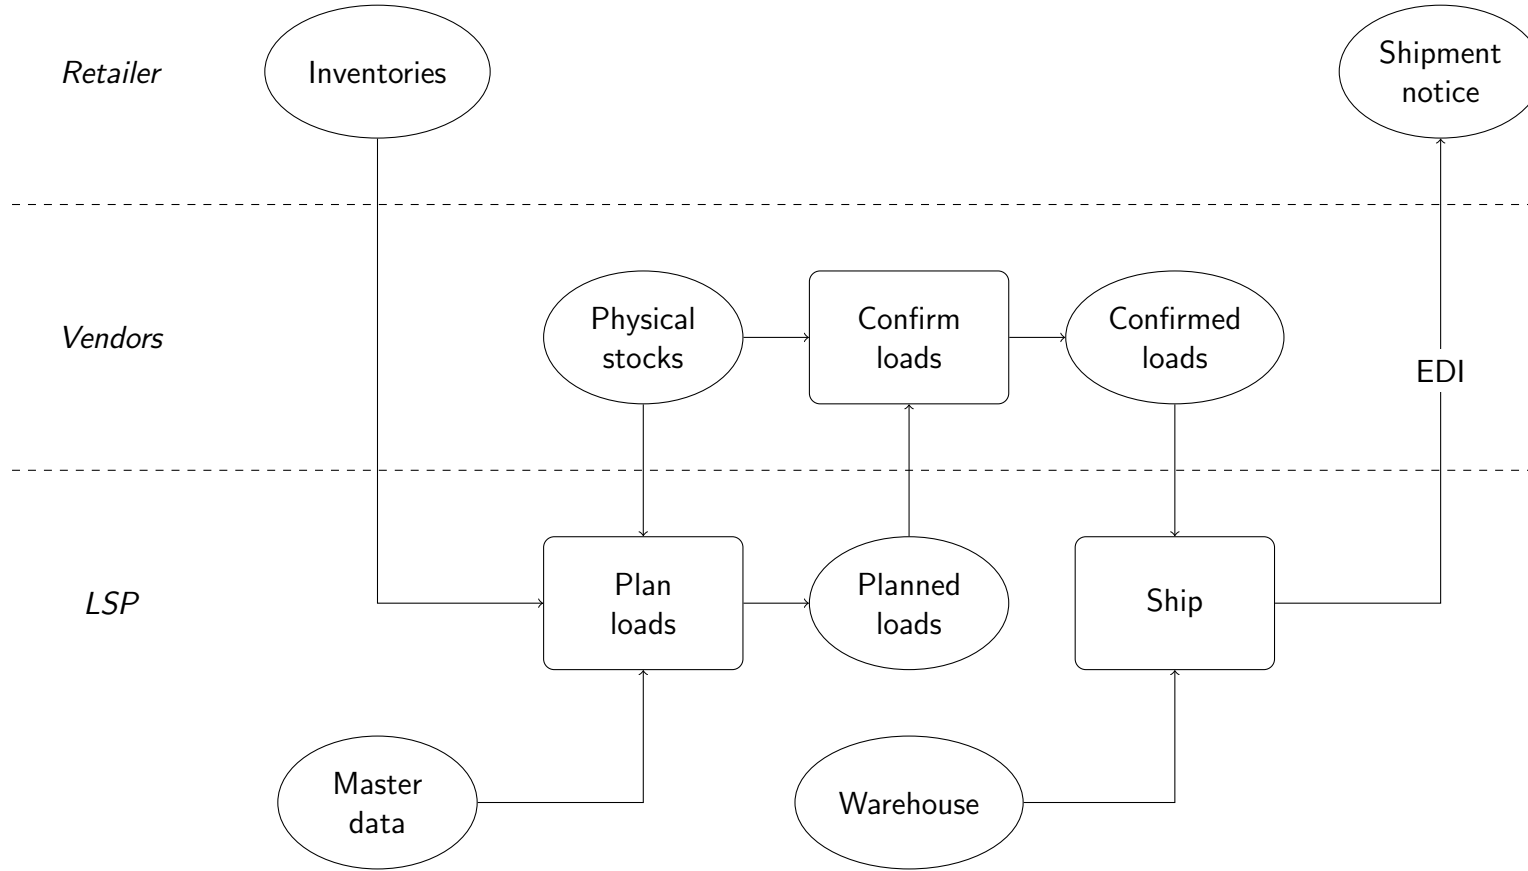


Figure 3.5 Shared VMI scenario process

3.5 Shared VMI scenario

Figure 3.5 shows the shared VMI scenario, which is totally different from the synchronization scenario. Its seemingly simple process actually contains the most complex decision: jointly managing the retailer inventory for all vendors' products while taking into account shipment efficiency and availability of physical stock.

IT landscape

Figure B.4 (page 55) shows the IT landscape for shared VMI. Note that naturally (but not necessarily), the LSP is also the orchestrator. The VMI software will most likely be made-to-measure, since multi-vendor VMI software with load building capabilities is not simply available on the market. The number of required IT interfaces is limited: only the retailer's and vendors' inventories need to be imported, with orders released as usual to the vendors' ERP systems. For the ██████ project in ██████, the IT investment was approximately €█████.

Organizational dynamics

Shared VMI is a highly 'integrated' solution and therefore requires deep trust between vendors and the retailer, an innovative LSP and a long-term focus. Again, perceived unfairness in gain sharing could be addressed by differentiating the pallet equivalent tariffs between vendors while measuring the collective drop size.

Operational performance

As shown in the process diagram (Figure 3.5), all possible sources of information are considered simultaneously in the shared VMI scenario: retailer inventories, vendor physical stocks, and product master data (for load building). This implies that, without digressing into operational policies (see Chapter 4), shared VMI can yield the highest operational performance.

While excluded from the process diagram for clarity, since the LSP orchestrates the shared VMI scenario, even the vehicle fleet availability (and diversity) can be leveraged in planning drops. From best practices in ██████, we learned that in addition to the initial IT investments, the annual running costs of shared VMI are approximately €█████ per retailer (meetings and fees).

Four

Coordination

In this chapter, we present our novel heuristic ‘push to fill’ policy for coordinating retailer inventory and transportation; after motivating why for our situation heuristics are preferred to optimization.

Heuristics vs. optimization

By ‘coordinating’ retailer inventory and (FTC) transportation, we imply optimizing some objective driven by order quantities and drop sizes. The supply chain-wide objective would be (minimizing) total long-term logistics costs, although these are currently incurred by different players.

The cost optimum may never be reached in practice due to a variety of reasons, the most evident being that no model completely captures reality. Within a model, defined by its assumptions, an *optimization* method performs an exhaustive and possibly implicit search to find the (model) optimum, which is not necessarily the optimum in practice. When such a procedure is infeasible due to constraints (e.g. on time), a *heuristic* method may be employed to produce a satisfactory solution.

In our model, we manage the inventory position of each product at the retailer by periodically deciding on the order quantity at the vendor(s), which in turn requires vehicles to ship the ordered products. The *commonality* of the resource ‘vehicle capacity’ makes this a particularly complex supply chain decision. Moreover, we have a (temporal) dependency between decisions now and in the future.

For multi-vendor shipment consolidation, a solution method should have the following properties:

Effective The method should lead to satisfactory performance for each of the players on the defined objective(s);

Tractable The method should be implementable with reasonable time and computing power resources;

Understandable The method should be understandable for the stakeholders, such that conflicts about outcomes are prevented.

We can categorize solution methods as *myopic* (‘short-sighted’) or *forward-looking* (long-sighted). Forward-looking methods take into account the aforementioned temporal dependency between decisions. In order to rigorously account for temporal dependency in a stochastic environment (i.e. future effects of current decisions), complex optimization

methods like *stochastic dynamic programming* should be employed. Unfortunately, such methods do not fulfil the tractability requirement. Therefore, we select a myopic solution method¹.

Then, within the class of myopic solution methods, we prefer a heuristic to optimization (e.g. a mixed integer program) because of the understandability criterium. In a coalition, decisions may not always be in the *direct* best interest of a player (i.e. pushing inventory to a retailer). When the decision maker can refer to simple rules that are agreed upon in advance (instead of a ‘black box’ optimization tool), conflicts can be prevented.

‘Push to fill’ policy

Now, we introduce our ‘push to fill’ policy for coordinating inventory and transportation. The retailer DC inventory is reviewed on a periodic basis, due to customer service and transportation schedules (i.e. *periodic review*). In addition, vendors mostly require retailers to order products in multiples of some order quantity or practical *batch size*, e.g. a case or a pallet [Larsen and Kiesmüller, 2007].

The designated *pull* policy for managing inventory under such constraints is an (R, s, nQ) policy [Hadley and Whitin, 1963], which implies that the inventory position IP is reviewed every R time periods, and upon review an integer ($n \geq 0$) multiple of batch size Q is ordered such that the inventory position becomes greater than or equal to s , the *reorder level*:

$$n = \min \{m \in \mathbb{N}_0 \mid IP + mQ \geq s\} = \max \left\{ \left\lceil \frac{s - IP}{Q} \right\rceil, 0 \right\}$$

However, pull policies are not concerned with efficiency in transportation. Therefore, we introduce the ‘push to fill’ policy. Upon review we perform Procedure 4.1².

By n_i^0 , we imply that $n_i^0 Q_i$ is the pull order quantity *before* coordinating inventory with transportation, driven purely by the current inventory position IP_i : the ‘pull batches’. The reorder level s_i is often defined by either contractual agreements (e.g. in case of VMI) or strict service level requirements. Therefore, the ‘coordinated’ order quantity $n_i Q_i$ should always be greater than or equal to $n_i^0 Q_i$.

Then, we calculate the number of drops necessary to ship the pull batches using the ‘conversion factor’ α_i which represents the capacity requirement per unit. Estimating the conversion factor can be straightforward (e.g. driven completely by weight in case of heavy products) or complex (e.g. when a combination of weight, volume, axle loading, stacking, etc. determines the capacity usage).

From the EOQ formula, we learn that the total costs increase rapidly for order quantities below the EOQ, as the setup cost is usually much higher than the holding costs. Similarly, we argue that the transportation costs are much more relevant than holding costs. Therefore, we suggest that whenever a vehicle is dispatched, it is completely filled by ‘pushing’ as many products downstream as necessary to fill the last vehicle. Pushing

¹Under certain assumptions on temporal dependency, myopic policies may even be provably optimal for the multi-period problem, see e.g. the base stock policy as a Newsvendor problem [Van der Vlist, 2007, p. 22].

²We index the decision variables and parameters for all products. See Appendix A for our definitions.

Procedure 4.1 Push to fill policy

```

1: for all  $i \in \mathcal{P}$  do
2:    $n_i^+ \leftarrow 0$  ▷ Initialize number of push batches
3: end for
4:
5: for all  $i \in \mathcal{P}$  do
6:    $n_i^0 \leftarrow \max \{ \lceil (s_i - IP_i) / Q_i \rceil, 0 \}$  ▷ Set number of pull batches
7: end for
8:
9:  $x \leftarrow \lceil (\sum_{i \in \mathcal{P}} \alpha_i n_i^0 Q_i) / C \rceil$  ▷ Fix number of drops
10:
11:  $\bar{\mathcal{P}}^+ \leftarrow \left\{ i \in \mathcal{P}^+ \mid \alpha_i Q_i \leq Cx - \sum_{j \in \mathcal{P}} \alpha_j (n_j^0 + n_j^+) Q_j \wedge \right.$ 
    $\left. IP_i + (n_i^0 + n_i^+ + 1) Q_i \leq s_i^+ \right\}$ 
12: while  $\bar{\mathcal{P}}^+ \neq \emptyset$  do ▷ Push still possible
13:    $i \leftarrow \arg \min_{j \in \bar{\mathcal{P}}^+} \left\{ (IP_j + (n_j^0 + n_j^+) Q_j - s_j) / E[D_j] \right\}$  ▷ Smallest run-out time
14:    $n_i^+ \leftarrow n_i^+ + 1$  ▷ Add push batch
15:    $\bar{\mathcal{P}}^+ \leftarrow \left\{ i \in \mathcal{P}^+ \mid \alpha_i Q_i \leq Cx - \sum_{j \in \mathcal{P}} \alpha_j (n_j^0 + n_j^+) Q_j \wedge \right.$ 
    $\left. IP_i + (n_i^0 + n_i^+ + 1) Q_i \leq s_i^+ \right\}$ 
16: end while
17:
18: for all  $i \in \mathcal{P}$  do
19:    $r_i \leftarrow (n_i^0 + n_i^+) Q_i$  ▷ Release order
20: end for

```

these products will lead to holding cost increases, but will also prevent costly vehicle dispatches in the future.

Since we suggest to completely fill vehicles, we can argue that the long term average transportation costs are optimized. However, arbitrarily choosing ‘push products’ could have negative effects, e.g. pushing slow movers which incur long term holding costs, or pushing products more urgently needed by other retailers.

Therefore, we introduce a *run-out time* (or coverage) based policy for ‘pushing’ products. First we introduce a set of ‘candidate’ push products $\mathcal{P}^+ \subseteq \mathcal{P}$ such that products $\mathcal{P} \setminus \mathcal{P}^+$ are never pushed³. The *iterative step* adds a ‘push batch’ of the product with the smallest expected run-out time of the ‘reorder level overshoot’ $IP_i + (n_i^0 + n_i^+) Q_i - s_i$, which still fits inside the vehicle (using $\alpha_i Q_i$ units of capacity) and will remain below its maximum inventory position for push s_i^+ ; and is repeated until *no* product fits⁴ or would remain below its maximum inventory position for push (i.e. $\mathcal{P}^+ = \emptyset$).

The reorder level s_i is often calculated as safety stock plus cycle stock, hence we deduct it from the projected inventory position $IP_i + (n_i^0 + n_i^+) Q_i$ to prevent always pushing products with lower safety stocks.

The push to fill policy relies on two assumptions:

³E.g. because of complex upstream production batching or retailer preferences.

⁴In practice, the push to fill policy can be implemented with iterative calls to a load builder to check the remaining capacity, instead of using a conversion factor.

Feasible capacity assignment When more than one drop is necessary ($x > 1$), we implicitly assume that while $\sum_{i \in \mathcal{P}} \alpha_i (n_i^0 + n_i^+) Q_i \leq Cx$, there exists a feasible assignment of these batches to vehicles such that no vehicle's individual capacity constraint C is violated.

Vendor inventory availability We ignore the availability of physical stock at the vendor DC. Of course, the policy can easily be extended by not ordering (pulling and pushing) out-of-stock products. Should such situations be frequent, then the reorder levels $\{s_i\}$ should be increased to maintain appropriate service levels at the retailer DC.

Five

Simulation

In this chapter, we show how the push to fill policy can be used to simulate different scenarios, which adjustments are necessary to implement the policy, and what the sequence of events in the simulation is.

5.1 Simulating different scenarios

Scenarios are characterized mainly by which player is coordinating inventory with transportation (see Chapter 3). For this purpose, we developed the push to fill policy which minimizes transportation costs subject to minimum inventory constraints (see Chapter 4).

We assume that orders at non-coordinated vendors are created using the simple (R_i, s_i, nQ_i) policy. Then, the push candidate set $\mathcal{P}^+ \subseteq \mathcal{P}$ can be defined to model the coordination of the push to fill policy. Table 5.1 shows an overview.

Table 5.1 Push to fill policy for different scenarios

Scenario	\mathcal{P}^+
Synchronization	\emptyset
Retailer orchestration	\mathcal{P}
Single VMI	\mathcal{P}_v with $v \in \mathcal{V}$ the VMI vendor
Shared VMI	\mathcal{P}

Synchronization is the basic scenario, which coordinates order schedules, but not orders, hence $\mathcal{P}^+ = \emptyset$. Then, for retailer orchestration and shared VMI, we deduce from the process diagrams in Chapter 3 that all products are coordinated (thus push candidates). For single VMI, the VMI vendor will decide on its order quantities (i.e. coordinate inventory with transportation) after the non-VMI vendors, but since we assume all non-VMI vendors use the (R_i, s_i, nQ_i) policy, we can simply simulate the push to fill policy for all vendors simultaneously while only pushing batches for the VMI vendor.

Since the feedback scenario is more pragmatic (and manually coordinated), it is complex to capture the coordination beyond synchronization in a simulation. We argue that the feedback loop will most likely be implemented with a single vendor, hence the savings would be smaller than but proportional to those in the single VMI scenario.

5.2 Implementing the push to fill policy

The push to fill policy in its purest form (see Chapter 4) needs to be adapted for implementation, both in practice and in simulation.

Periodic review and reorder level

In reality, the number of time periods between reviews R_i may not be constant. For example, the order schedule might repeat on a weekly basis: order on Monday, Wednesday, and Friday, while demand continues on Sunday. This implies that the reorder levels s_i should be adjusted to reflect the number of days until receipt of the next shipment.

We assume stationary demand at the retailer DC, which is backordered if necessary, characterized by an expected demand per time period¹ $E[D_i]$. We use a fixed lead time τ_L for all products (thus all vendors). We calculate the number of time periods until the next review² τ_R . Then, the *cycle stock* portion of reorder level s_i is equal to $(\tau_R + \tau_L) E[D_i]$, because the inventory position after the current order should cover all demand until the next order is *delivered*. Often, retailers set their *safety stock* in DCs to protect against unexpected upstream events (e.g. shipment delays or picking errors) *and* demand variability. A standard solution is incorporating a number of ‘safety time periods’ τ_S . Then, we have

$$s_i = (\tau_R + \tau_L + \tau_S) E[D_i] \quad \forall i \in \mathcal{P}.$$

For the maximum inventory position for push, we introduce τ^+ , the number of ‘push’ time periods allowed:

$$s_i^+ = s_i + \tau^+ E[D_i] \quad \forall i \in \mathcal{P}.$$

We note that since n_i^0 ensures that the projected inventory position is above s_i , no product for which $Q_i > \tau^+ E[D_i]$ will be pushed, i.e. for which the expected run-out time of a batch is larger than the number of allowed ‘push’ time periods³.

Sequence of events

In the simulation, we iterate over a range of time periods. We introduce the time index $t \in \mathcal{T}$ for all time-dependent parameters and decision variables, to structure the sequence of events, see Procedure 5.1. The procedure is rather self-explanatory. We suggest that $t^{\min} \geq \tau_M + 1$ to allow estimation of D_i . We initialize all inventory positions/levels at a multiple of the batch size Q_i . Note that $I_i^-(t) = \max\{IL_i^-(t), 0\}$ is the physical stock for which holding costs are incurred.

5.3 Evaluating results

We log the simulation results throughout (inventory and vehicle dispatch data), to compare different scenarios and vendor coalitions on the following KPIs:

Number of drops The total number of drops is equal to

$$X = \sum_{t \in (t^{\min}, \dots, t^{\max} - \tau_L)} x(t).$$

¹Often, the expected demand is forecasted.

²We assume that all products are reviewed simultaneously.

³ IP_i^- can become larger than s_i^+ due to pull orders, e.g. when batch size Q_i is large for a slow mover.

Procedure 5.1 Simulation

```

1: for all  $i \in \mathcal{P}$  do
2:    $E_{t^{\min}-1}[D_i] \leftarrow \left( \sum_{t'=t^{\min}-\tau_M}^{t^{\min}-1} D_i(t') \right) / \tau_M$ 
3:    $s_i(t^{\min}-1) \leftarrow (\tau_R(t) + \tau_L + \tau_S) E_{t^{\min}-1}[D_i]$ 
4:    $IP_i^-(t^{\min}) \leftarrow \lceil s_i(t^{\min}-1) / Q_i \rceil Q_i$  ▷ Initialize inventory position
5:    $IL_i^-(t^{\min}) \leftarrow IP_i^-(t^{\min})$ 
6: end for
7:
8: for  $t \leftarrow t^{\min}, t^{\max}$  do
9:   for all  $i \in \mathcal{P}$  do
10:     $IP_i(t) \leftarrow IP_i^-(t) - D_i(t)$  ▷ Satisfy demand
11:     $IL_i(t) \leftarrow IL_i^-(t) - D_i(t)$ 
12:     $B_i(t) \leftarrow \max \{ D_i(t) - I_i^-(t), 0 \}$  ▷ Log backorders
13:
14:     $E_t[D_i] \leftarrow \left( \sum_{t'=t-\tau_M+1}^t D_i(t') \right) / \tau_M$ 
15:     $s_i(t) \leftarrow (\tau_R(t) + \tau_L + \tau_S) E_t[D_i]$ 
16:     $s_i^+(t) \leftarrow s_i(t) + \tau^+ E_t[D_i]$ 
17:   end for
18:
19:   PUSH TO FILL ▷ Release orders using Procedure 4.1
20:
21:   for all  $i \in \mathcal{P}$  do
22:     $IP_i^-(t+1) \leftarrow IP_i(t) + r_i(t)$ 
23:     $IL_i^-(t+1) \leftarrow IL_i(t) + r_i(t - \tau_L)$  ▷ Receive shipments
24:     $I_i^-(t+1) \leftarrow \max \{ IL_i^-(t+1), 0 \}$  ▷ Physical stock carried over
25:   end for
26: end for

```

The total docking *setup* time at the retailer DC is proportional to X . Apart from comparing scenarios, it might also be interesting to quantify the sub-optimality of a scenario. In the case of drops, we can prove⁴ a theoretical lower bound:

$$X \geq \left\lceil \frac{\sum_{i \in \mathcal{P}} \alpha_i Q_i \left\lceil \frac{-IL_i^-(t^{\min}) + \sum_{t \in (t^{\min}, \dots, t^{\max} - \tau_L)} D_i(t)}{Q_i} \right\rceil}{C} \right\rceil.$$

Transportation costs We introduce a transportation cost function $f'(l)$ for the total drop size $l \in [0, \infty)$ per order, with

$$f'(l) = f \left(l - \left\lfloor \frac{l}{C} \right\rfloor \right) + f(C) \left\lfloor \frac{l}{C} \right\rfloor.$$

⁴*Simplified proof.* Because of our assumptions on $IL_i^-(t^{\min})$, we have

$$\sum_{t \in \mathcal{T}} n_i(t) \geq \left\lceil \frac{-IL_i^-(t^{\min}) + \sum_{t \in (t^{\min}, \dots, t^{\max})} D_i(t)}{Q_i} \right\rceil \quad \forall i \in \mathcal{P}. \quad \square$$

Then, the total transportation costs F are equal to

$$F = \sum_{t \in (t^{\min}, \dots, t^{\max} - \tau_L)} f' \left(\sum_{i \in \mathcal{P}} \alpha_i n_i(t) Q_i \right).$$

Again, we can calculate a theoretical lower bound for information on the maximum optimality gap in transportation costs:

$$F \geq f' \left(\sum_{i \in \mathcal{P}} \alpha_i Q_i \left[\frac{-IL_i^-(t^{\min}) + \sum_{t \in (t^{\min}, \dots, t^{\max} - \tau_L)} D_i(t)}{Q_i} \right] \right).$$

Physical stock DC inventory is often measured in pallet places. We introduce Q_i^0 for the number of units per pallet of product $i \in \mathcal{P}$. Then, we calculate the average number of *pallet places*⁵ at the retailer DC as

$$\frac{1}{t^{\max} - t^{\min} + 1} \sum_{i \in \mathcal{P}} \sum_{t \in (t^{\min}, \dots, t^{\max})} \left[\frac{I_i^-(t)}{Q_i^0} \right],$$

which is particularly interesting to investigate the inventory impact of pushing products. Alternatively, the average physical stock can be measured in days:

$$\frac{\sum_{i \in \mathcal{P}} \sum_{t \in (t^{\min}, \dots, t^{\max})} I_i^-(t)}{\sum_{i \in \mathcal{P}} \sum_{t \in (t^{\min}, \dots, t^{\max})} D_i(t)}.$$

Service level The *type 2* or β service level [Schneider, 1978, p. 1182] is defined as

$$\frac{\sum_{i \in \mathcal{P}} \sum_{t \in (t^{\min}, \dots, t^{\max})} B_i(t)}{\sum_{i \in \mathcal{P}} \sum_{t \in (t^{\min}, \dots, t^{\max})} D_i(t)}$$

and represents the fraction of demands satisfied directly (from on hand stock) in the retailer DC.

Based on these basic KPIs, we can make useful statements about all hypothesized effects of multi-vendor shipment consolidation (see Chapter 1). We are now ready to combine our knowledge about scenarios (and scenario selection), coordination, and simulation in a case study.

⁵Unlike in a vehicle, in the retailer DC any partial pallet occupies a full pallet place.

Scenario selection

Contingent upon the implementation of an incentivizing tariff structure, we argued that all vendors will benefit from synchronizing their shipments (no investment required). Therefore, we view synchronization as the basic scenario for multi-vendor shipment consolidation. From there, players should carefully consider the four other scenarios. We categorize the scenarios in a *scenario matrix* (see Figure 6.1) with two trade-offs: *inventory management* and *integration*.

Coordination of inventory and transportation is the responsibility of the inventory manager, which could be either the retailer or the vendor(s). Integration indicates the level of data transparency, interdependence of processes and size of investments.

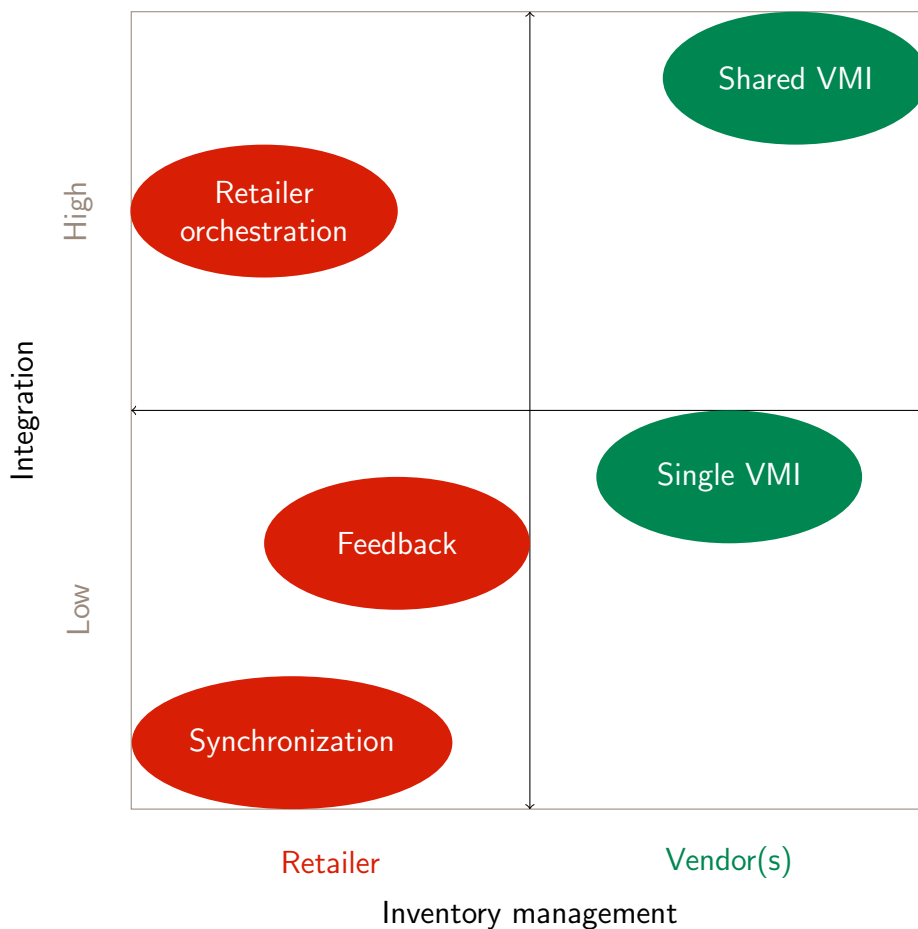


Figure 6.1 Scenario selection matrix

Scenario choice considerations can be structured by making these two trade-offs in the given order:

Inventory management Often, the retailer will have a strong preference for either RMI or VMI. Most likely, the incumbent strategy is aligned with this preference. However, when this preference is not decisive, we argue in favor of a vendor-managed scenario:

- From the process diagrams, we learn that in RMI scenarios, the vendor's physical stock is not taken into account in orders. Out of stock situations will therefore lead to unnecessary underutilization of vehicle capacity.
- Unless the vendors use factory gate pricing, they pay for FTC transportation. Therefore, vendors have a direct incentive to improve vehicle utilization and invest in innovations. The indirect benefits obtained by the retailer may not justify investments.
- Implementing (shared) VMI at the LSP will yield the additional benefit of incorporating the vehicle fleet availability in decisions.

Integration *Within* either RMI or VMI, we consider two scenarios: shared VMI is more integrated than single VMI and retailer coordination is more integrated than the feedback scenario. We note that decision makers should carefully investigate:

- Whether the higher investment for the integrated scenario yields significant additional benefits (using the simulation from Chapter 5);
- Whether the organizational landscape is strong enough to accommodate the integration;

before selecting the highly integrated scenario, and if this is not the case, whether the low-integration scenario still yields sufficient benefits on top of synchronization to warrant implementation.

We note that the above trade-offs are presented as if they pertain to the long-term optimal scenario. The scenario selection matrix can also be used to chart a path, e.g. transition to single VMI after experimenting with the feedback scenario.

Seven

Case study

As mentioned in Chapter 1, our *case class* is ‘multi-vendor shipment consolidation scenario selection problems’ and a *case* is defined by the combination of a multi-vendor DC and retailer DC. In this chapter, we address the scenario selection problem for the case ‘██████ to ██████’. The case study is structured using the *regulative cycle* (Figure 1.2).

7.1 Introduction

In this case, Heinz outsources warehousing and FTC transportation to ██████, which operates a multi-vendor DC in ██████ with co-located vendors ██████ and ██████. Multiple retailers were considered for the case study, and we selected ██████ (hereafter ‘██████’) because of their willingness to share the demand data necessary for our simulation (i.e. orders to the retailer DC). Specifically, we study ██████’s ██████ DC in ██████. This DC accounts for ██████% of the total volume originating from ██████¹. We shortly introduce the co-located vendors, LSP, and retailer in Appendix C.

The complete supply chain situation is shown in Figure 7.1, and we focus on improving the *LTL* transportation between the ██████ DC and ██████ DC.

¹██████ shared the Drops table (see Appendix E) for the complete retailer portfolio.




Figure 7.1 Supply chain situation




Figure 7.2 (Fictional) tariff structure

Table 7.1 Volume shipped from VDC by vendor (November 2012 to May 2013)

In this case, the working unit at the retailer DC is a ‘colli’, and the batch size Q_i may vary from pallet layer to full pallet (with Q_i^0 the number of products per pallet). Time periods are days, with the system operating seven days a week. The vehicles are trucks and the drop size is measured in ‘pallet equivalents’: (layers of) different products may be combined on a ‘pick pallet’. Hereafter, we will refer to pallet equivalents simply as ‘pallets’. A standard Dutch truck has a capacity of 26 pallet places. Pallets of products can be stacked, hence multiple pallets may occupy a single pallet place. Depending on the stackability of the vendors’ products, a truck may theoretically carry up to 52 pallets.

Figure 7.2 introduces the fictional² tariff structure $f(l)$ used throughout this case study. The tariffs are both representative of (i.e. proportional to) the invoices to vendors and the operational costs incurred by the LSP to make a drop of a certain drop size. Note that the cost savings potential for a drop size beyond ■ pallets is ■ (■%).

For reference, we show the volume distribution between the vendors in Table 7.1 and Figure 7.3 based on shipment data from the LSP. See Section E.1 for details on the LSP data and processing steps we performed. Clearly, Heinz ships the most pallets, which creates interesting organizational dynamics. The solid lines show the volume shipped from the multi-vendor DC in ■ to the retailer DC in ■. The dashed lines show the shipments from the ■ DC to the stores, based on retailer data (see Section E.2,

²The fictional tariff structure was provided by ■ and used for all vendors to prevent anti-trust.




Figure 7.3 Volume shipped by vendor and period (November 2012 to May 2013)

which also explains the slight mismatch for █████). There is a noticeable bullwhip effect [Lee et al., 1997], which could be caused by the buffering effect of physical stock at the retailer DC and/or batching in the orders to the (multi-)vendor DC.

7.2 Problem

The orders by █████ to the co-located vendors are *not coordinated* with transportation, which is a missed opportunity for multi-vendor shipment consolidation. Thus, the players face the *scenario selection problem*. The scenario selection problem is particularly relevant for the █████ DC, because many (Dutch) retailers are served from this DC and more vendors may join in the future, which implies that the selected scenario might be rolled out and/or scaled up.

7.3 Diagnosis

To get a firmer grasp of the problem and create a basis for our redesign, we diagnose the current drop size, costs, and docking times.

Drop size

From Figure 7.4 we learn that the drop sizes at █████'s █████ DC are far from optimal: while a full truck contains *at least* 26 pallets (due to the possibility of stacking),




Figure 7.4 Histogram of drop size (November 2012 to May 2013)

the current average drop size is █████ pallets³. We note that these are *drop sizes*, which are not necessarily the truck loads from █████ to █████. The LSP normally consolidates these ‘drops’ with other nearby drops. But these drop sizes *do* form the basis for invoices.

In total, █████ drops were made at █████, whereas █████ drops (a █████% reduction) would have sufficed for shipping the total volume (assuming 26 pallets per drop). As mentioned above, this estimate is conservative due to the possibility of stacking. The number of drops drives docking setup time (see Figure 7.6 below), volume-independent handling costs, etc. (see Figure 1.3). █████ estimates that the volume-independent handling costs for a single drop are €████ [Van Moorsel, 2009].

Table 7.2 Volume and drop size by address in pallets (November 2012 to May 2013)

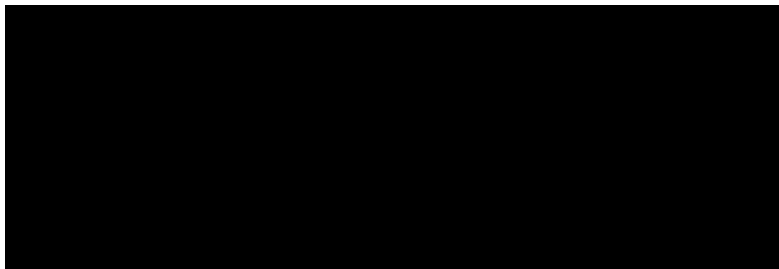
Upon further inspection, we learn that █████’s █████ DC actually has docks

³Although the drop size seems to be increasing over time, illustrated by a linear coefficient of █████ for the date sequence, this effect is non-significant ($p = 0.03$).

on different addresses, called ‘ship to’ addresses⁴, see Table 7.2, which require separate drops (and have separate drop sizes and corresponding invoices). Some addresses are characterized by extremely low average drop sizes, especially given the large number of visits (i.e. volume divided by average drop size). Moreover, the different vendors use different addresses, which prevents multi-vendor shipment consolidation.

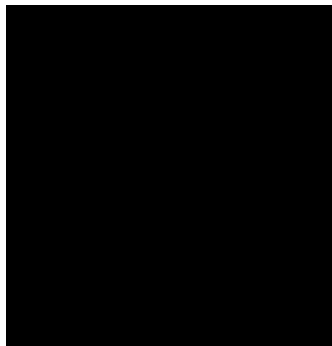
It is in the vendors’ best interest to have a single ‘ship to’ address at a retailer DC, since this is directly reflected in the drop size and therefore in the transportation costs. Hence, we recommend that █████ is contacted to sort out these inconsistencies. In our redesign, we assume that the █████ DC has a single ship to address, which is both optimal (for the vendors) and necessary given our lack of data on which products are shipped to which address.

Table 7.3 Volume and drop size by weekday in pallets (November 2012 to May 2013)



In addition, the weekdays for shipments are not coordinated. Although █████’ schedule is to ship to █████ on Tuesday and Friday (the lead time of $\tau_L = 2$ days implies ordering on Sunday and Wednesday), many shipments take place on Monday and Thursday, see Table 7.3. Most likely, these are retailer orders prioritized by the customer service departments of the vendors. There is also a relation between vendors and weekdays, with Friday used predominantly for █████ (especially given its overall █% share from Table 7.1).

Table 7.4 Demand by weekday in pallets (November 2012 to May 2013)



It is not only relevant to synchronize the ordering schedule between vendors, but also with the downstream demand at the retailer DC: adhering to the just-in-time (JIT) principle to minimize physical stock. In Table 7.4, we show the demand at the retailer DC (shipments to stores) per weekday in pallets. With a total demand of █████ pallets in seven months, roughly █████ pallets should be shipped to █████ per week. This implies

⁴A single retailer may have multiple ship to’s in the vendors’ systems, e.g. due to the internal lay-out of the retailer DC.



Figure 7.5 Costs for distribution of one FTL between two drops

that with a vehicle capacity of 26 pallets, a schedule of at most ■ inventory reviews per week should be implemented.

Costs

The total costs⁵ for the drops (November 2012 to May 2013) at ■ were €■. The ‘conservative⁶’ optimum for the total volume is €■, indicating a ■% cost savings potential. The average costs per pallet equivalent were €■.

To understand the effects of volume distribution between drops, we investigate the tariff structure. We observe that the cost increase for adding a pallet equivalent (or ‘slope’ of the tariff structure) is incidentally non-decreasing, see the costs for the first drop in Figure 7.5. When 26 pallets are distributed between two drops, of course it is optimal to ship one FTL. However, the most expensive choice is counterintuitively to make two drops of 13 pallet equivalents. The transportation costs are minimized for (close-to-)FTL drops with incidental small drops if necessary, which disfavors continuously shipping in LTL at a higher frequency.

Docking time

It is particularly relevant to consolidate the ship to addresses and synchronize the schedules in light of docking times. Although the ship to addresses for ■ may be geographically nearby; for each drop, the truck has to ‘setup’ before unloading. From Figure 7.6, we can conclude that per drop the average docking *setup* time is ■ minutes⁷, with an additional ■ minutes unloading time per pallet equivalent (both the intercept and slope coefficient are significant at the $p < 0.001$ level).

⁵Using the tariff structure from Figure 7.2.

⁶We use ‘conservative’ from here onwards for assuming one pallet equivalent per pallet place.

⁷Setup time means all volume independent docking time, e.g. arrival, departure, etc.




Figure 7.6 Diagnosis of docking time (November 2012 to May 2013)

7.4 Design

Based on the diagnosis, we believe that there is need for a *redesign* of the ordering and shipment process. For the redesign, we use the ‘design knowledge’ [Van Aken, 2005, p. 30] on multi-vendor shipment consolidation scenarios developed in Chapter 3. We follow the scenario selection process outlined in Chapter 6. ██████ does not have a decisive preference for RMI or VMI, hence all scenarios are considered.

Simulation

We implemented the simulation (see Procedure 5.1 on page 26) in Microsoft Excel. ██████ has shared the historical demand⁸ $D_i(t)$ to the retailer DC in ██████. In addition, we have acquired the necessary ‘master’ data (the set \mathcal{P} with parameters α_i and Q_i for each $i \in \mathcal{P}$). See Section E.2 for more details on the retailer data and the processing steps we performed. Appendix F summarizes all underlying assumptions for reference.

Most notably, we calculate the conversion factor α_i , the capacity requirement per unit, in *(ISO) pallet equivalents*⁹. We use $C = 26$, i.e. a truck can carry 26 pallet equivalents on its 26 pallet places. This is both conservative (a truck can realistically carry up to ██████ pallets for the current vendors due to stacking) and progressive (we assume ‘perfect picking’¹⁰).

⁸I.e. cumulative daily shipments to the ██████ stores.

⁹E.g. when a ISO pallet contains 120 units of product $i \in \mathcal{P}$, then $\alpha_i = 1/120$.

¹⁰E.g. 120 different products with $\alpha_i = 1/120$ fit on a single pallet place.

Table 7.5 Simulated schedule for both frequencies

We use a $\tau_M = \blacksquare$ day moving average to forecast the demand per time period. We analyze $t^{\max} - t^{\min} + 1 = 243 - 32 + 1 = 212$ time periods. The lead time is $\tau_L = \blacksquare$ days and we incorporate $\tau_S = \blacksquare$ safety days (also used by Heinz for VMI). We use $\tau^+ = \blacksquare$, i.e. the vendors are allowed to push up to \blacksquare days above the reorder level s_i to fill trucks. This gives us a best case for pushing, as this equals the average number of days on hand and allows doubling physical stock for a product.

Initially, we simulate two scenarios for two frequencies (see below). Comparing synchronization ($\mathcal{P}^+ = \emptyset$) and retailer orchestration/shared VMI ($\mathcal{P}^+ = \mathcal{P}$) provides insight into the ‘best case for coordination’. The impact of coordination was moderate, so we additionally simulated the medium-integrated single VMI scenario with Heinz as the best VMI vendor (see the volume distribution in Table 7.1).

Schedule

The order (and thereby shipment) schedule is an important driver of supply chain performance. The arrival of a shipment and time until the next shipment arrives are the only levers available to influence the physical stock (demand is viewed as an exogenous stochastic process). Moreover, demand during the period between orders determines ‘pull’ drop sizes.

\blacksquare assumes that there is no demand on \blacksquare . Hence, we modify the calculation of the reorder level $s_i(t)$ in Procedure 5.1:

$$s_i(t) = \left(\tau_R(t) + \tau_L + \tau_S - \mathbf{1}_{(w(t) \bmod 7) + \tau_R(t) + \tau_L \geq 7} \right) \mathbb{E}_t[D_i] \quad \forall i \in \mathcal{P},$$

thus if the period until the next shipment includes Sunday, then a day is deducted from the required coverage through the indicator function $\mathbf{1}$.

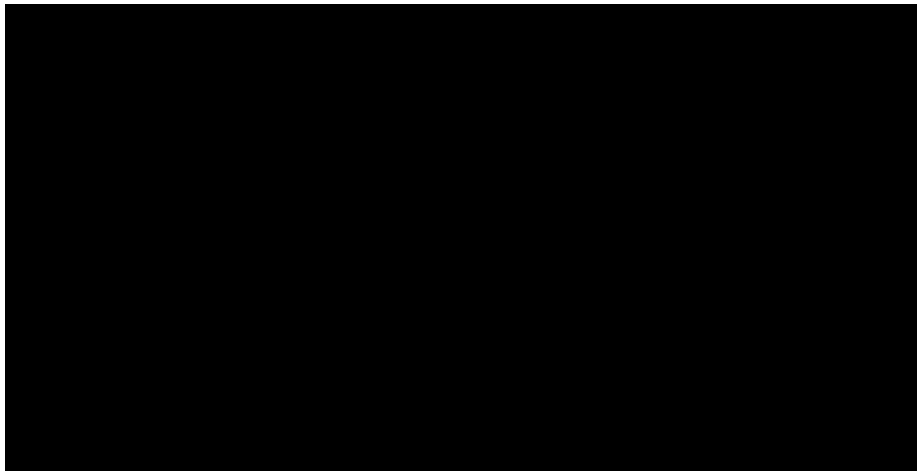
\blacksquare currently ships to \blacksquare on \blacksquare and \blacksquare , but from Table 7.4 we learn that demand on \blacksquare and \blacksquare is negligible. Moreover, \blacksquare is unable to receive shipments on Saturday and Sunday, thus physical stock must be held over the weekend for covering Monday and shipping on Monday is beneficial to prevent also holding stock for Tuesday over the weekend. When assuming a fixed leadtime (see Appendix F) ordering on Thursday or Friday is irrational, since the lead time will increase and later orders will be shipped simultaneously but based on new information.

We simulate two frequencies: ordering twice per week (‘2’) and three times per week (‘3’). Further frequency increases required unevenly spaced schedules and yielded too low drop sizes while stock decreases were negligible. Given the above considerations, the schedules in Table 7.5 performed best, with evenly spaced drop sizes. The schedule of frequency 2 is perfectly just-in-time for the busiest days Tuesday and Friday (see Table 7.4).

Results

Below, we analyze the drops and inventory during the simulation. All results are ‘conservative’ as introduced earlier, based on the assumption of one pallet equivalent per pallet place. We refer to Figure 1.3 for the hypothesized effects of multi-vendor shipment consolidation, which are driven by drops and inventory. In Appendix G, the results are extensively validated.

Table 7.6 Simulation results for drops (November 2012 to May 2013)



Drops

Table 7.6 shows the drop-related KPIs resulting from the simulation of our scenarios and frequencies (lower bounds calculated as in Section 5.3). When we take a closer look at the transportation costs, it turns out that merely by:

- Synchronizing (and keeping to) the schedule;
- Modifying the schedule (to ordering on [REDACTED] and [REDACTED]);
- Consolidating the ship to addresses;

the as-is costs of €[REDACTED] can be reduced by [REDACTED]% to €[REDACTED]. In addition, the number of drops can be reduced by [REDACTED]% to [REDACTED]. This corresponds to €[REDACTED] in savings on handling in the retailer DC.

Implementing shared VMI at frequency 2 can yield [REDACTED]% additional (compared to synchronization) savings on the total transportation costs and further reduce the number of drops to [REDACTED] ([REDACTED]%, €[REDACTED] in handling) by shipping only FTLs. Our ‘push to fill’ policy is therefore successful. Single VMI is just slightly more constrained in pushing (see explanation below) and can yield additional transportation cost savings of [REDACTED]%.

Figure 7.7 shows a histogram of the drop size for each scenario and frequency. Synchronization (naturally) yields the best drop sizes at a frequency of 2 instead of 3: [REDACTED] drops ([REDACTED]%) are already FTL *without* coordination. At frequency 3, the additional inventory for push of $\tau^+ = [REDACTED]$ days slightly constrains the fill procedure of shared VMI. This effect is more pronounced for single VMI due to the smaller set of push products. When all products are at their maximum inventory position for push, filling becomes impossible, leading to incidental LTL drops.




Figure 7.7 Simulation results for drops (November 2012 to May 2013)

Inventory

Table 7.7 Simulation results for inventory (November 2012 to May 2013)

Table 7.7 shows the inventory-related KPIs resulting from the simulation. Interestingly, we disprove a hypothesized effect from Figure 1.3: the service level¹¹ decreases for synchronization when we increase the frequency to 3. The only possible explanation is the schedule shown in Table 7.5. The corresponding physical stock decrease is almost negligible (■%).

¹¹We note that in reality, service levels are much higher due to the use of more advanced forecasts to determine orders, e.g. promotions, product introductions, weekday and seasonal patterns, etc., which is not the focus of this study. The observed service level is therefore a lower bound.




Figure 7.8 Simulation results for inventory (November 2012 to May 2013)

Shared VMI at frequency 2 strongly increases the service level by ■ percentage points through a ■% physical stock increase. This confirms that our ‘run-out time’ based push to fill policy successfully selects push products. Single VMI leads to a slightly lower physical stock than shared VMI, since the set of push products is limited¹². Increasing the frequency counterintuitively increases the physical stock for shared VMI, since every truck at the higher frequency needs to be filled (see Figure 7.8 above).

Hence, also taking drop-related KPIs into account, increasing to frequency 3 does *not* make sense for the ■ case.

Business case

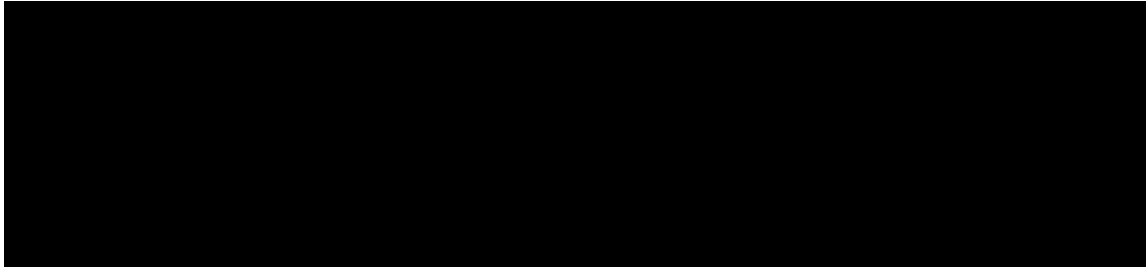
As mentioned in Chapter 3, synchronization is a ‘zero investment alternative’. Since the synchronization scenario outperforms the as-is situation on all KPIs, the more integrated scenarios need to be evaluated against synchronization instead of the as-is performance. For reference, the current annual transportation spend under synchronization for ■ would be approximately €■.

The business case is based on estimates (ballpark figures) obtained from interviews and simulation at frequency 2 under the assumptions in Appendix F. The investment is highly specific to the incumbent IT landscape and its standardization. The simulated

¹²Note that the average number of pallet places for single VMI at frequency 2 is higher than for shared VMI, while the average number of days is lower. This is caused by the ‘ceiling’ effect of pallet places: as soon as one unit is stored, a pallet place is necessary.

savings are based on demand realizations for seven months and extrapolated to expected yearly savings by applying a factor 12/7. Table 7.8 summarizes the business case for █████, with transportation cost savings obtained by the vendors and handling/holding cost savings obtained by the retailer¹³.

Table 7.8 Business case



Retailer orchestration Since retailer orchestration has no running costs, it seems a better scenario than shared VMI. However, by pushing products the savings on handling costs and increased holding costs approximately cancel out, and there is no return on the investment for the retailer.

Single VMI The single VMI scenario still yields transportation cost savings, against a lower investment than shared VMI, depending on the necessity of load building software. There may be multiple reasons to implement VMI, hence it is unclear whether the running costs should be included in the business case¹⁴. The transportation cost savings of █% alone are not sufficient for implementation.

Shared VMI We observe that the shared VMI scenario does not even cover its running costs with transportation savings, and hence should not be implemented for the █████ case.

Generalization to network

Approximately █████ pallets are shipped to █████ each year by █████. This represents only █% of the total volume originating from █████ (around █████ pallets annually). It is therefore interesting to generalize the business case for multi-vendor shipment consolidation to the total volume.

The average drop size for all retailers is █████ pallets. Figure 7.9 shows that a moderate number of drops is already (close to) FTL, but a large number of drops is for *less than █ pallets*. Figure 7.10 visualizes the retailer ship to addresses in The Netherlands replenished from █████. Many retailers fall into the same category as █████: less than █████ pallets shipped per year with an average drop size of less than █████ pallets.

However, █████ indicates that █████'s situation of multiple ship to addresses *and* non-synchronized schedules is quite unique (today). Figure 7.9 does not look much better than Figure 7.4, but some small drops are a commercial necessity (e.g. to retailers with a low volume but strategic relevance).

¹³Note that the holding cost savings are negative due to pushing.

¹⁴E.g. Heinz is already using VMI at █████.




Figure 7.9 Histogram of drop size for all retailers (November 2012 to May 2013)

We therefore generalize the savings from Table 7.8 to the complete customer portfolio using a factor of 25%, to remain realistic given ██████'s situation. This implies portfolio-wide synchronization could collectively save the vendors €█████ annually on the transportation costs of €█████ (█%). If we assume handling a drop costs each retailer €█ and the portfolio-wide number of drops (█████) can be reduced by █%, this amounts to €█████ saved on retailer handling costs alone annually.

Similar to the ██████ project in ██████, if implemented, shared VMI should be a concept at the vendor DC rolled out to as many retailers as possible to leverage the investment. When we assume that shared VMI can lead to █% savings for every retailer, which is *highly* speculative, then the (retailer-specific) running costs of €█████ would only be covered for retailers with an annual (synchronized) spend of more than €█████.

There aren't many Dutch retailers with such large volumes, and on top of that the savings should generate a return on the initial investment. We conclude that rolling out shared VMI from the ██████ DC is not suitable for the Dutch market due to the typically low transportation costs (and moderately convex shape of the tariff structure).




Figure 7.10 Map of ship to addresses of all retailers

7.5 Recommendations

Based on the business case, our recommendation for the ██████████ case is simple:

Implement synchronization The players should synchronize their ship to addresses¹⁵ and order schedules *as soon as possible*. This will *annually* save up to €████ on transportation costs (████%) and up to €████ on handling costs in the retailer DC (████%). All players benefit and no investment is required.

Moreover, ████████ could investigate simple ways to incorporate the vendors' collocation in ████████ into ordering decisions to prevent the most costly drops in the future (slight retailer orchestration).

In addition, our recommendations for the ████████ network (vendors and LSP) are:

Improve tariff structure In this case study, we calculated the transportation costs based on a fictional tariff structure. However, in practice the vendors pay different tariffs and are invoiced for their *individual* drop size. Such a tariff structure does not incentivize multi-vendor shipment consolidation. ████████ has indicated that redesigning the tariff structure has a high priority (see Appendix D).

¹⁵ ████████ was already considering to consolidate its ship to addresses during a planned ████████ label implementation.

Experiment with single VMI Multi-vendor shipment consolidation using single VMI is promising, but not as a stand-alone initiative. However, Heinz is discussing VMI with another retailer. This is an excellent opportunity to experiment with pushing products to fill vehicles given the other vendors' drop sizes.

Roll out to all retailers We recommend that the vendors roll out synchronization to all retailers. This could annually save up to €■■■■ on transportation costs and up to €■■■■ on handling costs, without any investment. Moreover, the network should approach large retailers with a business case for retailer orchestration, showing that they can additionally reduce the number of drops by up to ■■%.

Eight

Reflection

We complete the reflective cycle (Figure 1.2) by reflecting on the case study and improving our design knowledge. Three interesting observations from the case study are discussed below.

8.1 Increasing the shipping frequency

Increasing the shipping frequency is cited as one of the two main benefits of multi-vendor shipment consolidation (see Figure 1.3). But the shipping frequency is often studied in *continuous time*, e.g. Çetinkaya and Lee [2000] optimize the shipping frequency for VMI in conjunction with upstream batch sizes, and Coppens [2012, p. 26] analyzes the expected frequency increase for multi-vendor shipment consolidation as the inverse of the accumulation time of an FTL.

In reality, the food retail supply chain operates in *discrete time*: transportation and demand are regulated by a schedule. The shipping frequency can therefore only be influenced through the schedule, which is constrained by practicalities (e.g. no goods receipt during the weekend, the same schedule every week, etc.). Implementing a theoretical frequency in an operational schedule may therefore lead to performance loss.

Certainly, when multiple smaller vendors co-locate, they can improve their shipping frequency. But as the total volume increases, it becomes harder to reach the theoretical frequency. More integrated scenarios than synchronization merely optimize the transportation costs given a schedule, but by pushing products partially cancel the stock reduction effect of a higher frequency. Once synchronized, these scenarios do not lead to a higher shipping frequency, they merely enable it from a transportation costs perspective.

Based on this observation, we expand the design knowledge for multi-vendor shipment consolidation: co-located vendors can improve their *schedule*, and thereby their frequency. In a case study, the focus should therefore be on designing a better schedule instead of always increasing the frequency.

8.2 Tariff structure

One of the most striking observations in the case study is that while multi-vendor shipment consolidation scenarios can reduce the number of drops by more than ■%, the corresponding transportation cost reduction is approximately ■%. Multi-vendor shipment consolidation is very effective, but we must not forget that logistics service providers have essentially already provided vendors with access to shipment consolidation for decades.

Just like consumers can send parcels to the other side of the world for say €■, vendors can send a pallet from ■■■■■ to ■■■■■ for just €■■■■. LSPs have created their whole business model around drop consolidation, and cost competitiveness (due to the moderate entry barriers of the industry) has driven them to perfection.

When vendors would not outsource transportation (or pay the full transportation costs for every drop), the tariff structure would be extremely convex, and synchronization could save them €■■■■ annually for ■■■■■ alone and up to €■■■■ for the entire retailer portfolio¹. Multi-vendor shipment consolidation is therefore much more urgent in other countries, where the tariff structure is more convex due to a lack of consolidation opportunities for the LSP, or where the transportation costs are much higher due to longer distances [e.g. in the US, see Hanson Logistics, 2013].

This reinforces our design knowledge by viewing multi-vendor shipment consolidation as a truly multi-objective initiative (see Figure 1.3), not just as a way to reduce transportation costs. The reduction of the number of drops directly reduces the docking setup times (non-value added waiting time for the vehicle) and handling costs; and helps reduce carbon emissions and congestion.

8.3 Effectiveness of coordination

In Chapter 4 we have developed a new ‘push to fill’ policy to optimize the drop size. From Figure 7.7, we learn that at a higher frequency, scenarios based on push to fill coordination start to underperform. In the case study, the coefficient of variation of demand is ■■%, which is quite high for a downstream echelon like the retailer DC. The stochasticity will at almost every inventory review require some batches (pull orders), even though the run-out time based push policy attempts to prevent pull orders at the next review. Due to our assumption that pull batches are strictly necessary, there is always a drop and hence a vehicle to be filled. At some point, all candidate push products will be at their maximum inventory position due to push, leading to LTL drops. We therefore believe that push to fill coordination performs best for a case with relatively low variation of demand, to truly prevent drops by pushing products with a low run-out time.

With a strong schedule (see above), synchronization will already lead to acceptable drop sizes. As mentioned above, push coordination can further reduce the number of drops by more than ■■%, but this is not proportionally reflected in transportation cost savings. Hence, the drop size optimization offered by coordination will also be more beneficial for more *convex* and/or higher tariffs.

¹Assuming €■■ per drop, with €■■ deducted from the FTL costs for volume-dependent costs.

We conclude this thesis by answering the research questions and providing future research directions.

Research questions

What are the (dis)advantages of multi-vendor shipment consolidation?

From the start of the project, we have had hypotheses about the (dis)advantages of multi-vendor shipment consolidation. The issue tree in Figure 1.3 has been adapted numerous times after interviews with experts, discussion sessions and the study of best practices, before finally being subjected to quantitative analysis in our case study.

The simulation results confirmed almost all our hypotheses, but we must nuance the advantage of a higher shipping frequency [observed in Coppens, 2012]: only through better schedule design will vendors be able to fully reap the benefits of a higher shipping frequency. We thereby identify a gap between the continuous time based shipment consolidation literature and discrete time practice, which should be investigated in future research (see below).

What are feasible multi-vendor shipment consolidation scenarios?

Through multiple iterations with stakeholders, we have identified five multi-vendor shipment consolidation scenarios: synchronization, feedback, retailer orchestration, single VMI, and shared VMI. In Chapter 3, we present a process (re)design, insight into the organizational dynamics, and a blueprint for the IT landscape for each scenario.

How to coordinate multi-product inventory and shipment decisions?

In Chapter 3, we learned that improvement beyond synchronization requires coordination of order quantities with drop sizes. We assume that through correct parameterization, each inventory review leads to strictly necessary ‘pull batches’ using an (R, s, nQ) policy. Then, given the capacity requirement, we calculate the number of drops and fill the vehicles with ‘push batches’. This novel ‘push to fill’ policy successfully optimizes the drop size (see Table 7.6), but also leads to an increase in downstream physical stock (see Table 7.7).

What is the impact of different multi-vendor shipment consolidation scenarios?

In addition to describing the ‘qualitative’ impact of each scenario in Chapter 3, we developed and implemented a simulation procedure (Procedure 5.1) to quantify the operational performance of each scenario. We summarized each scenario’s performance in a comprehensive business case (see Table 7.8).

How to select a multi-vendor shipment consolidation scenario?

To guide players in selecting the appropriate scenario for each case, we presented a scenario selection framework (Figure 6.1) which summarizes the two key trade-offs: inventory management (RMI vs. VMI) and integration. Players should only select a more integrated scenario when the organizational landscape is ready and the additional savings generate an acceptable return on investment.

For our case study, synchronization is necessary but sufficient, due to the low convexity of the (Dutch) tariff structure, which is based on the LSP’s ability to already consolidate drops. The additional optimization of the drop size offered by coordination, which could further reduce the number of drops by more than ■%, does not generate a positive return of investment.

In summary, our main contributions are the identification and introduction of five multi-vendor shipment consolidation scenarios, the development of a new ‘push to fill’ policy which optimizes drop sizes given inventory requirements, and a scenario selection framework.

Future research directions

Future research directions for Heinz resulting from this project are simple: ■ (has indicated it) needs a new tariff structure for multi-vendor shipment consolidation, subject to the requirements in Appendix D. As prerequisite for multi-vendor shipment consolidation from the ■ DC, this research direction should be prioritized.

We identify three gaps in academic literature which deserve future research:

Analytical model The push to fill policy effectively optimizes transportation costs while improving the service level. Future research could build upon our findings by attempting to analytically model the pull and push quantities, to allow for a faster and more reliable assessment of operational performance than through simulation.

Discrete time order schedules In reality, order/shipment processes are dictated by schedules, which yield many operational advantages, but constrain inventory optimization and may lead to non-stationary demand between inventory reviews. Discrete time inventory management literature does not discuss schedule design, hence future research should include evaluating and optimizing order schedules.

Tariff structures While widespread in practice, the optimization of *multi-product* orders given a tariff structure for transportation is not addressed in literature. The order quantity results of Swenseth and Godfrey [2002] should be extended to the multi-product variable order quantity case.

Moreover, gain sharing has long been a discipline within cooperative game theory [e.g. Shapley, 1953], but (again) its practical application in tariff structures for multi-customer services of a provider has remained unaddressed. A post-hoc distribution of savings is not always desirable (see Appendix D), but the design of ex-ante mechanisms requires additional research.

A

Definitions

A.1 Sets

\mathcal{P} Set of products

$\mathcal{P}^+ \subseteq \mathcal{P}$ Set of candidate push products

$\mathcal{P}_v = \{i \in \mathcal{P} \mid v_i = v\}$ Set of products sold by vendor $v \in \mathcal{V}$

$\mathcal{T} = (1, \dots, t^{\max})$ Set of time periods

\mathcal{V} Set of vendors

A.2 Parameters

α_i Conversion factor of product $i \in \mathcal{P}$ (in ISO pallets per unit)

B_i Backordered demand during time period (in units)

C Vehicle capacity

h_i Holding cost rate per time period of product $i \in \mathcal{P}$

$I_i^- = \max\{IL_i^-, 0\}$ Physical stock at the retailer DC of product $i \in \mathcal{P}$ *before* satisfying demand, i.e. carried over from previous time period (in units)

IL_i Inventory level at the retailer DC of product $i \in \mathcal{P}$ *after* satisfying demand (in units)

IL_i^- Inventory level at the retailer DC of product $i \in \mathcal{P}$ *before* satisfying demand (in units)

IP_i Inventory position at the retailer DC of product $i \in \mathcal{P}$ *after* satisfying demand (in units)

IP_i^- Inventory position at the retailer DC of product $i \in \mathcal{P}$ *before* satisfying demand (in units)

Q_i Batch size (in units)

Q_i^0 Pallet size (in units)

s_i Reorder level of product $i \in \mathcal{P}$ (minimum inventory position in units)

s_i^+ Maximum inventory position for push of product $i \in \mathcal{P}$ (in units)

τ_L Lead time (in time periods)

τ_M Time periods for moving average

$\tau_R(t)$ Time periods until next review on time $t \in \mathcal{T}$

τ_S Time periods for safety stock

τ^+ Time periods for ‘push’ inventory

$v_i \in \mathcal{V}$ Vendor of product $i \in \mathcal{P}$

A.3 Functions

$f(l)$ Transportation costs for a drop size $l \in [0, C]$ in a single vehicle

$\mathbf{1}$ Indicator function, returns 1 if condition is true, 0 if condition is false

$w(t) \in (1, \dots, 7)$ Weekday of time $t \in \mathcal{T}$ (week starts on Monday)

A.4 Stochastic variables

D_i Demand for product $i \in \mathcal{P}$ during a time period (in units)

A.5 Decision variables

n_i Number of batches ordered for product $i \in \mathcal{P}$

n_i^0 Number of pull batches ordered for product $i \in \mathcal{P}$

n_i^+ Number of push batches ordered for product $i \in \mathcal{P}$

r_i Order released for product $i \in \mathcal{P}$ (in units)

x Number of trucks dispatched

B

Scenario IT landscapes

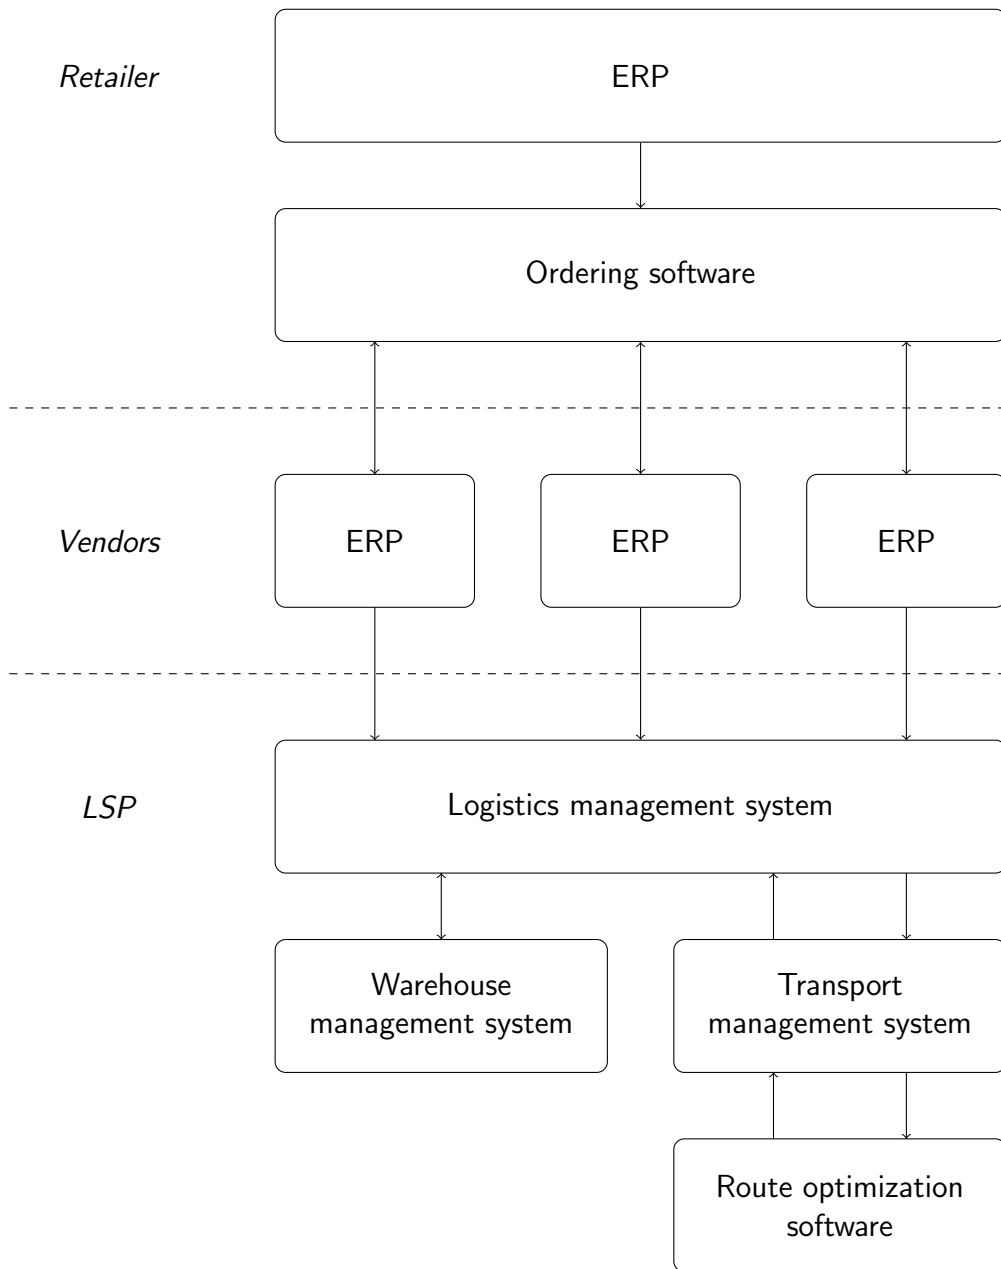


Figure B.1 Synchronization scenario IT landscape

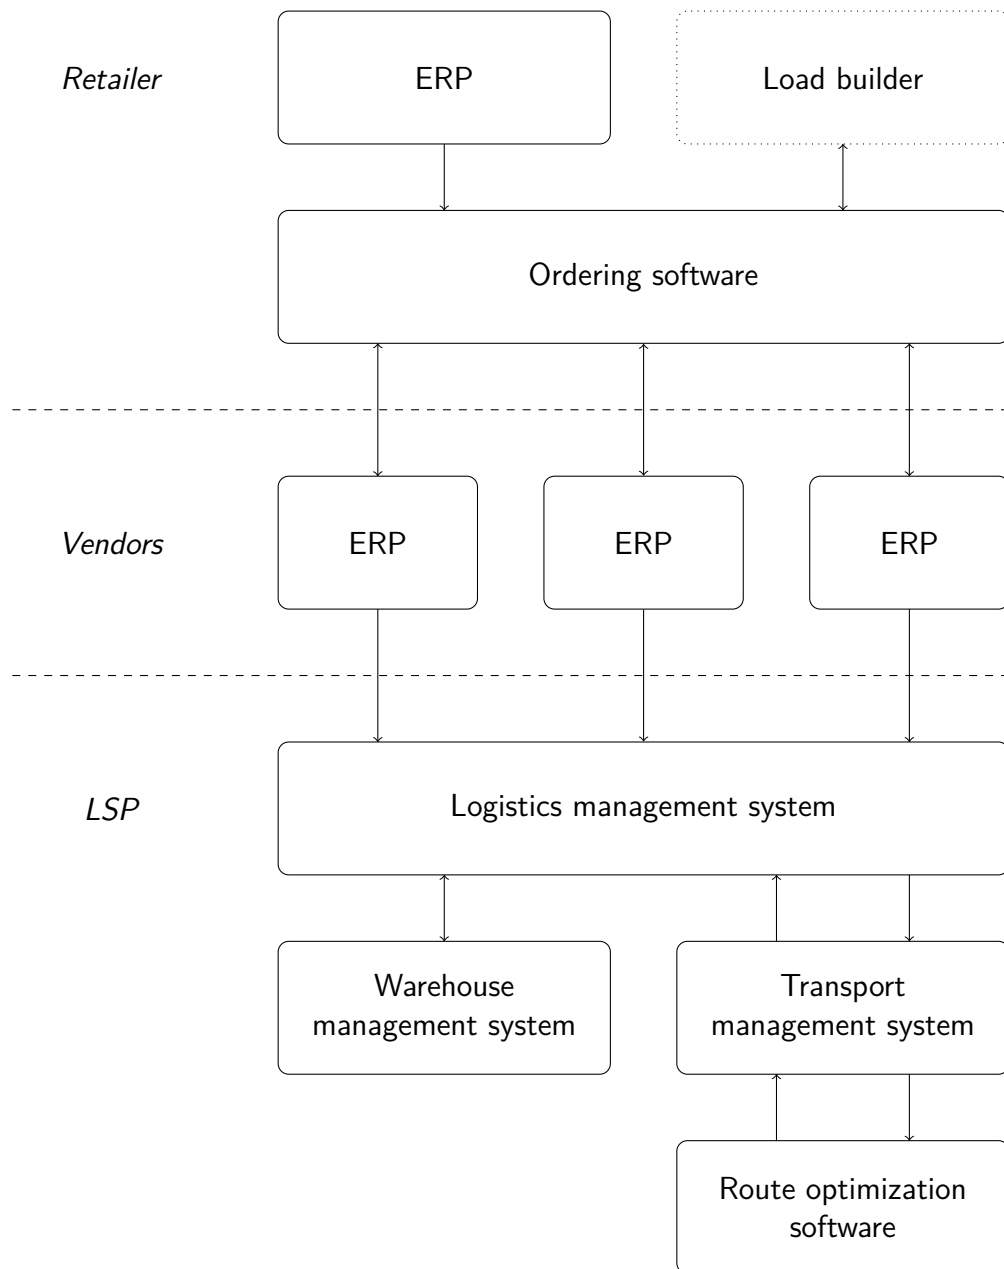


Figure B.2 Retailer orchestration scenario IT landscape

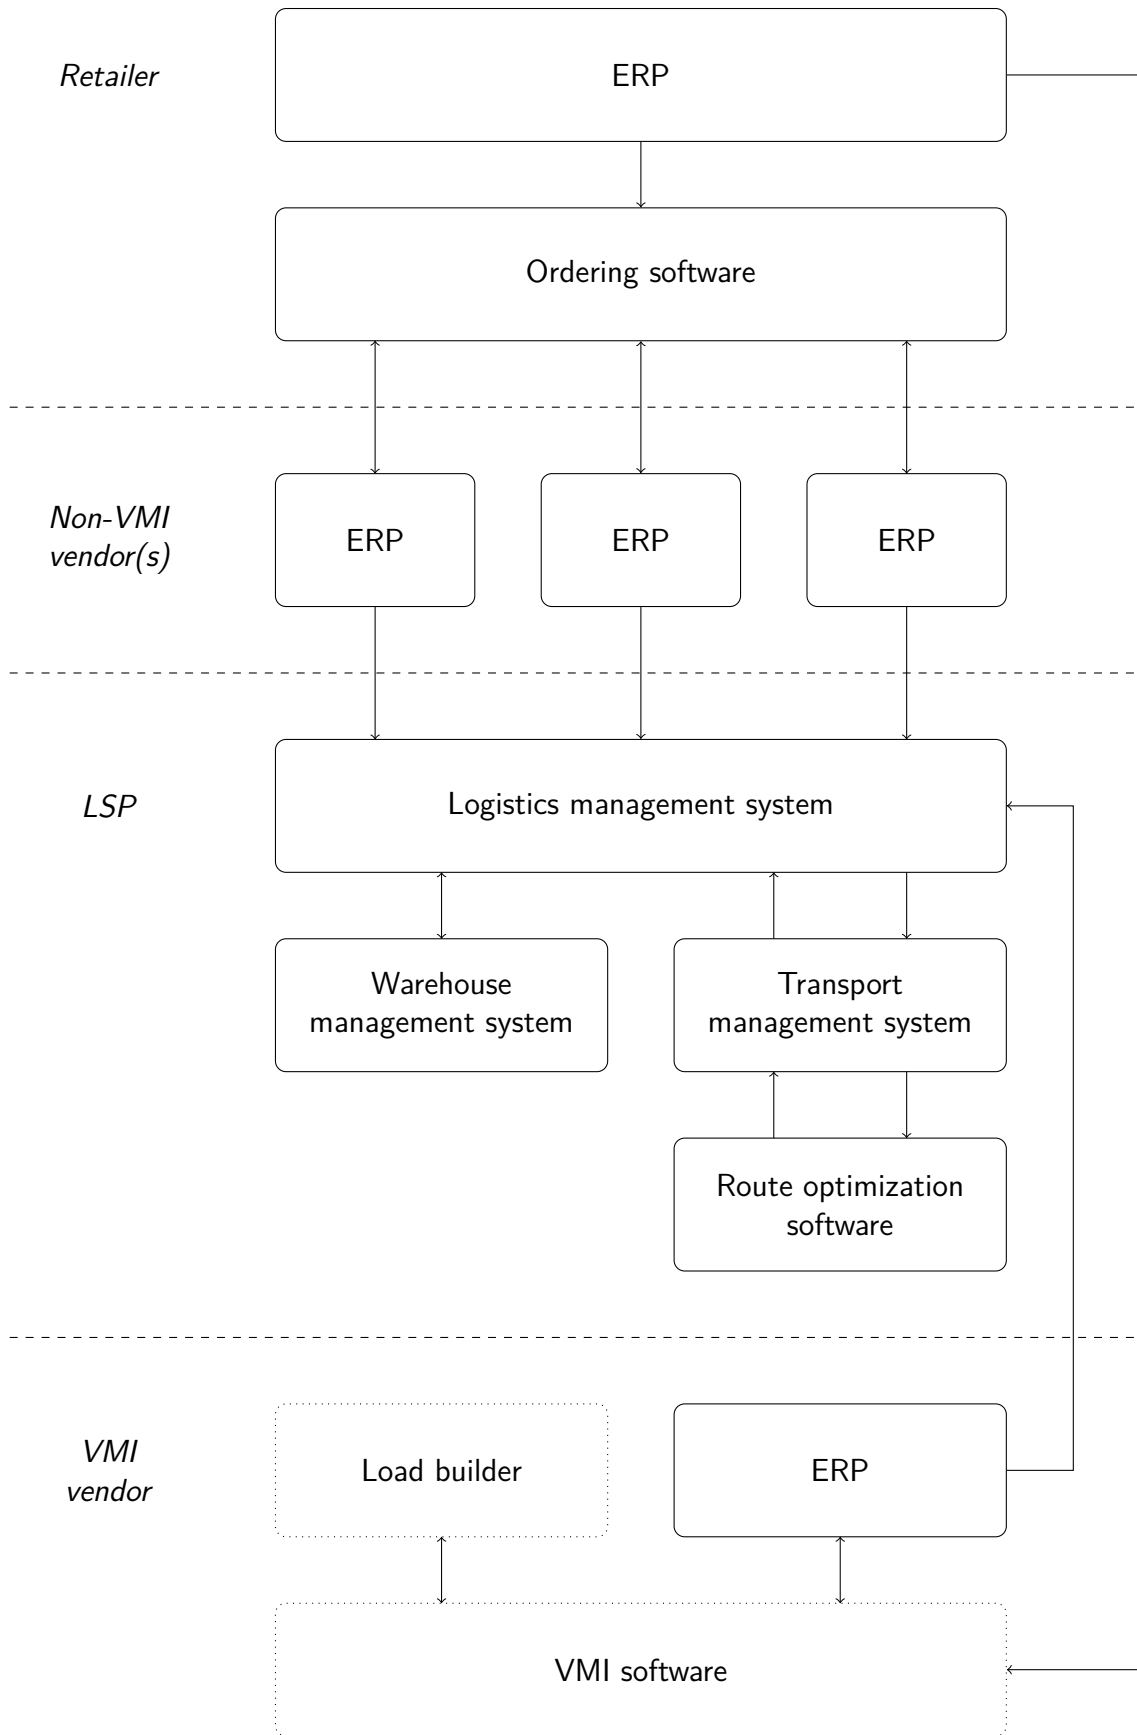


Figure B.3 Single VMI scenario IT landscape

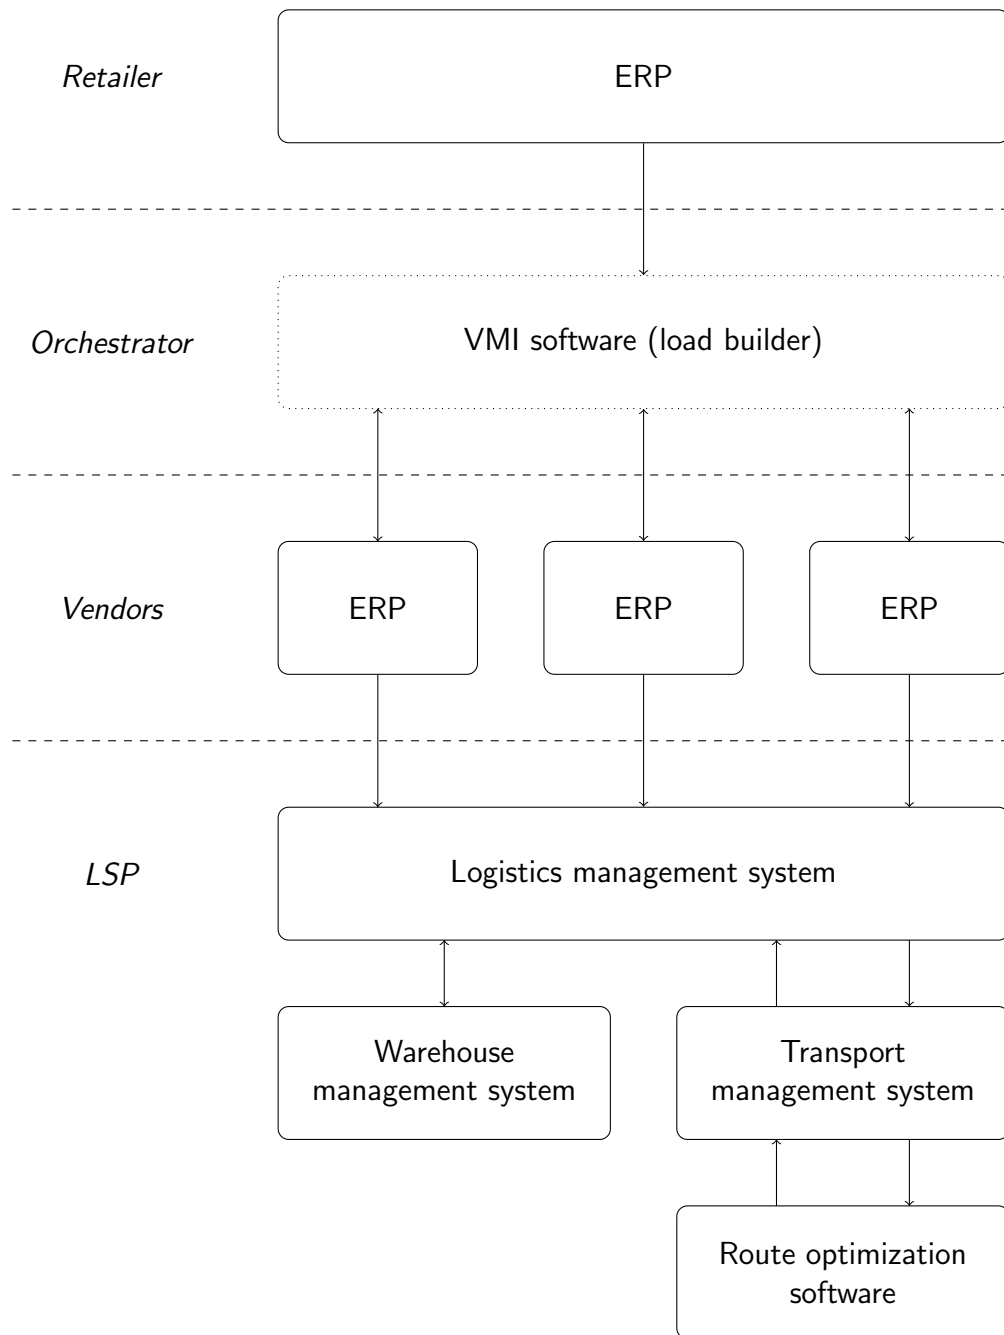


Figure B.4 Shared VMI scenario IT landscape

C

Company introductions

D

Gain sharing

As mentioned in Chapter 7, ██████ currently invoices vendors based on their *individual* drop size. This tariff structure does not incentivize multi-vendor shipment consolidation (e.g. increasing the frequency decreases the individual drop size, which increases the tariff). Therefore, the tariff structure needs to be redesigned.

Moreover, the transportation cost savings obtained through multi-vendor shipment consolidation should be shared between the players ('gain sharing'). The LSP can share in the gains by setting the margin between costs and tariffs. Then, for the vendors, the new tariff structure should satisfy the following properties:

Fair Gains should be proportional to contributions, because otherwise 'in the long run, some participants will inevitably become frustrated since their true share in the group's success is undervalued' [Cruijssen et al., 2007, p. 32].

Incentive No player should be harmed by the collaboration (e.g. pay more due to a higher frequency). This corresponds to the game theoretic concept of *stability* [Slikker, 2012].

Pro-active From interviews with experts, we learn that 'reactive' (post-hoc) gain sharing poses operational risk (pay more in anticipation of refunds), which prevents participation.

The redesign of ██████' tariff structure is out of scope for this project, but we present two best practices (and their limitations):

Fixed rate Using contractual agreements on the (average) shared drop size, the LSP can safely invoice vendors using a fixed rate (possibly differentiating the rate between the vendors). The limitation is that, apart from contractual agreements, the fixed rate can lead to indifference for the drop size at the vendors' customer service departments.

Shared drop size ██████ currently experiments with invoicing ██████ and ██████ using a tariff based on their shared drop size. For example, the costs for shipping 13 pallets in a shared FTL are 13 times the vendor's pallet equivalent tariff for an FTL.

The LSP's revenue will strongly decrease when the current tariffs are used based on the shared drop size, but a tariff increase will frustrate vendors who can individually ship FTL drops. Another limitation is that the *assignment* of more than an FTL to different drops on a given day influences the shared drop size (and hence tariffs paid by the vendors).

E

Data

In this appendix, we provide more details about the data we received and the processing steps we performed before analysis and simulation.

E.1 LSP data

For diagnosis and validation, ██████ has provided us with the `Drops` table, containing all *actual* shipments between November 1st 2012 and May 31st 2013 from the vendor DC in ██████ to the retailer DC in ██████. The `Drops` table contains the following columns (omitting some irrelevant columns):

████ Drop number

████ Date

████ Origin

████ Destination

████ Destination address

████ Drop size (in pallet equivalents)

████ Drop size (in pallet equivalents)

████ Drop size (in pallet equivalents)

████ Drop size (in pallet equivalents)

████ Drop size (in pallet equivalents)

████ Drop size (in pallet equivalents)

████ Shared drop size (in pallet equivalents)

████ Tardiness (early/on time/late)

████ Docking time (in minutes)

We filter on the destination and only keep the rows for which ██████ is equal to ██████ or ██████, thereby removing Christmas shipments and other irregular volumes. In addition, we remove two dedicated FTL shipments for ██████ (see Section E.2 below).

E.2 Retailer data

From [REDACTED], we received a `Shipments` table with shipments from the [REDACTED] DC in [REDACTED] to the stores, containing the following columns:

[REDACTED] Article number
 [REDACTED] Units per consumer unit
 [REDACTED] Type of consumer unit
 [REDACTED] Article name
 [REDACTED] Vendor number
 [REDACTED] Vendor name
 [REDACTED] Date in YYYYMMDD format
 [REDACTED] Units shipped
 [REDACTED] Units ordered
 [REDACTED] Units ordered but not shipped (not backordered)
 [REDACTED] European article number
 [REDACTED] Units per layer
 [REDACTED] Units per pallet
 [REDACTED] Consumer units per unit
 [REDACTED] Area within [REDACTED] DC
 [REDACTED] Minimum order quantity (in units)
 [REDACTED] Incremental order quantity (in units)
 [REDACTED] Start date of promotion
 [REDACTED] End date of promotion
 [REDACTED] Consumer price (per unit)

The `Shipments` table contains the shipments between October 1st 2012 and May 31st 2013. We define additional columns:

Vendor Simplified vendor name: Heinz (for [REDACTED]), [REDACTED] (for [REDACTED], [REDACTED], and [REDACTED]), or [REDACTED] (for [REDACTED])

Date Computer readable date

t Time period, calculated as `Date - 2012-09-31` in days¹

¹ $t^{\min} = 32$ is November 1st, 2012.

alpha_i Conversion factor, calculated as $1/\blacksquare$ for Heinz and $0.8/\blacksquare$ for \blacksquare and \blacksquare^2

h_i Holding costs³, calculated as \blacksquare

We filter the **Shipments** table and remove the rows with DC equal to \blacksquare or \blacksquare (cooled products, not originating from \blacksquare). In addition, \blacksquare uses direct shipment (from factory to retailer DC) for some of the products in the assortment of its subsidiary \blacksquare . \blacksquare provided a list of products shipped from \blacksquare and five products with mostly direct shipments were removed, due to the inability to match \blacksquare 's shipments to stores with \blacksquare shipments.

To obtain the ‘master’ data, we remove duplicates on \blacksquare from the **Shipments** table and map each row to a product $i \in \mathcal{P}$:

$i \leftarrow \blacksquare$

$Q_i \leftarrow \blacksquare$

$Q_i^0 \leftarrow \blacksquare$

$\alpha_i \leftarrow \text{alpha_i}$

$h_i \leftarrow \text{h_i}$

$v_i \leftarrow \blacksquare$

In total, the vendors sell $|\mathcal{P}| = \blacksquare$ products to \blacksquare .

To obtain the ‘demand’ data necessary for simulation, we map each row of the **Shipments** table to a product $i \in \mathcal{P}$ and time period $t \in \mathcal{T}$:

$i \leftarrow \blacksquare$

$t \leftarrow \blacksquare$

$D_i(t) \leftarrow \blacksquare$

The total of \blacksquare is larger than the actual demand of the stores due to \blacksquare ; but the \blacksquare are caused by historical out-of-stocks in the RDC, which are endogenous to the simulation. These \blacksquare cannot simply be deducted from future orders to reconstruct the ‘true’ demand pattern, hence we must assume that the demand is equal to the historical shipments \blacksquare . Also, when orders are indeed modified after \blacksquare , this ‘new insight’ more accurately reflects the demand pattern.

² \blacksquare and \blacksquare ship on \blacksquare pallets, while Heinz ships on ISO pallets, in which the truck capacity is measured. Hence \blacksquare uses a factor 0.8 for \blacksquare and \blacksquare to convert to ISO pallets.

³ \blacksquare internally uses $\blacksquare\%$ of the consumer price as annual inventory costs.

Assumptions

Below, we summarize and motivate all assumptions underlying the simulation for the ██████████ case study:

Backordering at retailer DC In reality, ██████████ does not backorder unsatisfied store demand at the retailer DC (miss orders). New orders are created upon the store’s next inventory review with new information. However, due to our inability to include this new information (e.g. consumer demand, inventory write-offs, etc.), we must assume that store orders to the retailer DC are backordered. This prevents long-term underestimation of the demand (see Section E.2, where we decide to use actual shipments and exclude historical miss orders to prevent overestimation).

Feasible capacity assignment When more than one drop is necessary ($x > 1$), we implicitly assume that while $\sum_{i \in \mathcal{P}} \alpha_i (n_i^0 + n_i^+) Q_i \leq Cx$, there exists a feasible assignment of these batches to vehicles such that no vehicle’s individual capacity constraint C is violated. The impact of this assumption is moderate because of the conservative assumption of 26 pallet equivalents per vehicle.

Fixed lead time We assume a fixed lead time of █ days with the system operating 24/7. In reality, orders are subject to approval by the vendors’ customer service desks and picking by the logistics service provider, which both do not continue during the weekend. E.g. it is still possible to deliver on Monday, but those orders should be released on ██████████ instead of ██████████ (which is used in the simulation). This effectively introduces a variable weekday-dependent lead time.

The physical stock (inventory level) will always be driven by the time until the next shipment arrives, so in this sense the simulation is realistic. However, orders are based on the inventory position, thus will be subject to more demand variability and based on less information when released earlier. This effect is out of scope due for the sake of simplicity.

Schedule compliance The simulation does not perfectly capture reality including human behavior in case of disruptions. Priority shipments outside the schedule are considered a ‘fact of life’ at Heinz, while the retailer inventory is not even reviewed outside of the schedule in the simulation. Simulation will show that keeping to the schedule can yield significant savings. While these savings are ‘best case’ due to perfect schedule compliance, the necessity of priority shipments should not be taken for granted.

Stationary demand Since the focus of this project is not forecasting, in our forecasts we assume that demand is stationary (i.e. no weekday patterns, seasonal effects, and

promotions), while this is not the case. The resulting service levels are acceptable but should be interpreted with caution.

Vendor inventory availability We ignore the availability of physical stock at the vendor DC, i.e. retailer demand is always satisfied. Of course, the push to fill policy can easily be extended by not ordering (pulling and pushing) out-of-stock products. Should such situations be frequent, then the reorder levels s_i should be increased to maintain appropriate service levels at the retailer DC. (Moreover, historical vendor DC inventory data was both unavailable and would be partially endogenous to the simulated decisions.)

Vehicle capacity We assume that the vehicles carry ■ pallet equivalents on their 26 pallet places. This is both conservative (due to stacking of up to ■ pallets of the current vendors in a single truck) and progressive (as it assumes perfect stacking to complete pallets of different products).

Schruben [1980] introduces a framework for validation of simulation models (Figure G.1). We subsequently qualify, verify, and validate our model.




Figure G.1 Validation of simulation models [Schruben, 1980]

G.1 Model qualification

Model qualification determines the adequacy of the conceptual model. [REDACTED] currently uses an (R, s, nQ) policy to release orders to the vendors, hence our push to fill policy (Procedure 4.1) with $\mathcal{P}^+ = \emptyset$ adequately represents the order process in reality. We model the capacity requirement of orders using a *conversion factor* α_i . Both Heinz and [REDACTED] internally use a conversion factor to estimate the number of pallets, which is calculated in the same way. The sequence of events (Procedure 5.1) was created in extensive dialogue with [REDACTED]. The conceptual model is therefore an adequate representation of reality.

G.2 Model verification

Next, we verify the computer model, which is an implementation of the conceptual model (Procedure 4.1 and Procedure 5.1) in Microsoft Excel. We run the simple case where $\mathcal{P} = \{1\}$, $\mathcal{P}^+ = \emptyset$, $Q_1 = 1$, $\alpha_1 = 1$, $\mathcal{T} = (1, \dots, 5)$, $\tau_L = \tau_S = \tau_M = 1$, $t^{\min} = 2$, and $D_1(t) = 1$ for all $t \in \mathcal{T}$. Orders are released (i.e. inventory is reviewed) every time period. Table G.1 verifies the simulation. Note that $s_i(t) = (\tau_R + \tau_L + \tau_S) E[D_i] = 3$, for all $t \in \mathcal{T}$.

Table G.1 Verification of simulation in Microsoft Excel ($\mathcal{P}^+ = \emptyset$)

t	$D_1(t)$	$IP_1(t)$	$IL_1(t)$	$r_1(t)$	x
$t^{\min} = 2$	1	2	2	1	1
3	1	2	1	1	1
4	1	2	1	1	1
$t^{\max} = 5$	1	2	1	1	1

Next, Table G.2 verifies the ‘push to fill’ implementation with $\mathcal{P}^+ = \mathcal{P}$ and $C = 2$. As expected, once a batch is pulled, an additional batch is pushed (filling the vehicle), which prevents an order in the next period.

Table G.2 Verification of ‘push to fill’ in Microsoft Excel ($\mathcal{P}^+ = \emptyset$)

t	$D_1(t)$	$IP_1(t)$	$IL_1(t)$	$r_1(t)$	x
$t^{\min} = 2$	1	2	2	2	1
3	1	3	1	0	0
4	1	2	2	2	1
$t^{\max} = 5$	1	3	1	0	0

G.3 Model validation

Finally, we validate the simulation results of the computer model with reality. There are two main sources of validity: statistical power and face validity.

Statistical power

We simulate on the historical demand realizations for a period of only 212 days (approximately 60 inventory reviews). Hence, we need to make sure that the results have sufficient statistical power.

The statistical power of the simulation pertains to robustness to the stochastic demand, as all other factors are (assumed to be) deterministic. The main focus of the simulation is the drop size, which is driven by the demand between inventory reviews. For synchronization, the demand between inventory reviews even directly determines the drop size.

Since the vendors’ assortments jointly contain ████ products, there is a strong ‘pseudo statistical’ effect: the demand is stochastic, but its amplitude (or standard deviation) is




Figure G.2 Demand for different instances

dampened by the consolidation to total drop size¹. The variance of results will therefore be much smaller than the variance of demand, which increases the statistical power of the drop size results.

Unfortunately, we cannot simply fit distributions to the stochastic demand to generate instances, due to seasonality, weekday patterns, product introductions, promotions, etc. Therefore, we add ‘random noise’ to the demand data provided by ██████:

$$D_i(t) \leftarrow \lfloor (0.5 + \text{RAND}()) D_i(t) + 0.5 \rfloor \quad \forall i \in \mathcal{P}, t \in \mathcal{T}$$

with $\text{RAND}()$ the pseudo-random number generator for a uniform distribution on $[0, 1]$ available in Microsoft Excel². This procedure inflates the variance of the stochastic demand variables while preserving the mean, and thereby allows us to analyze the sensitivity of our results.

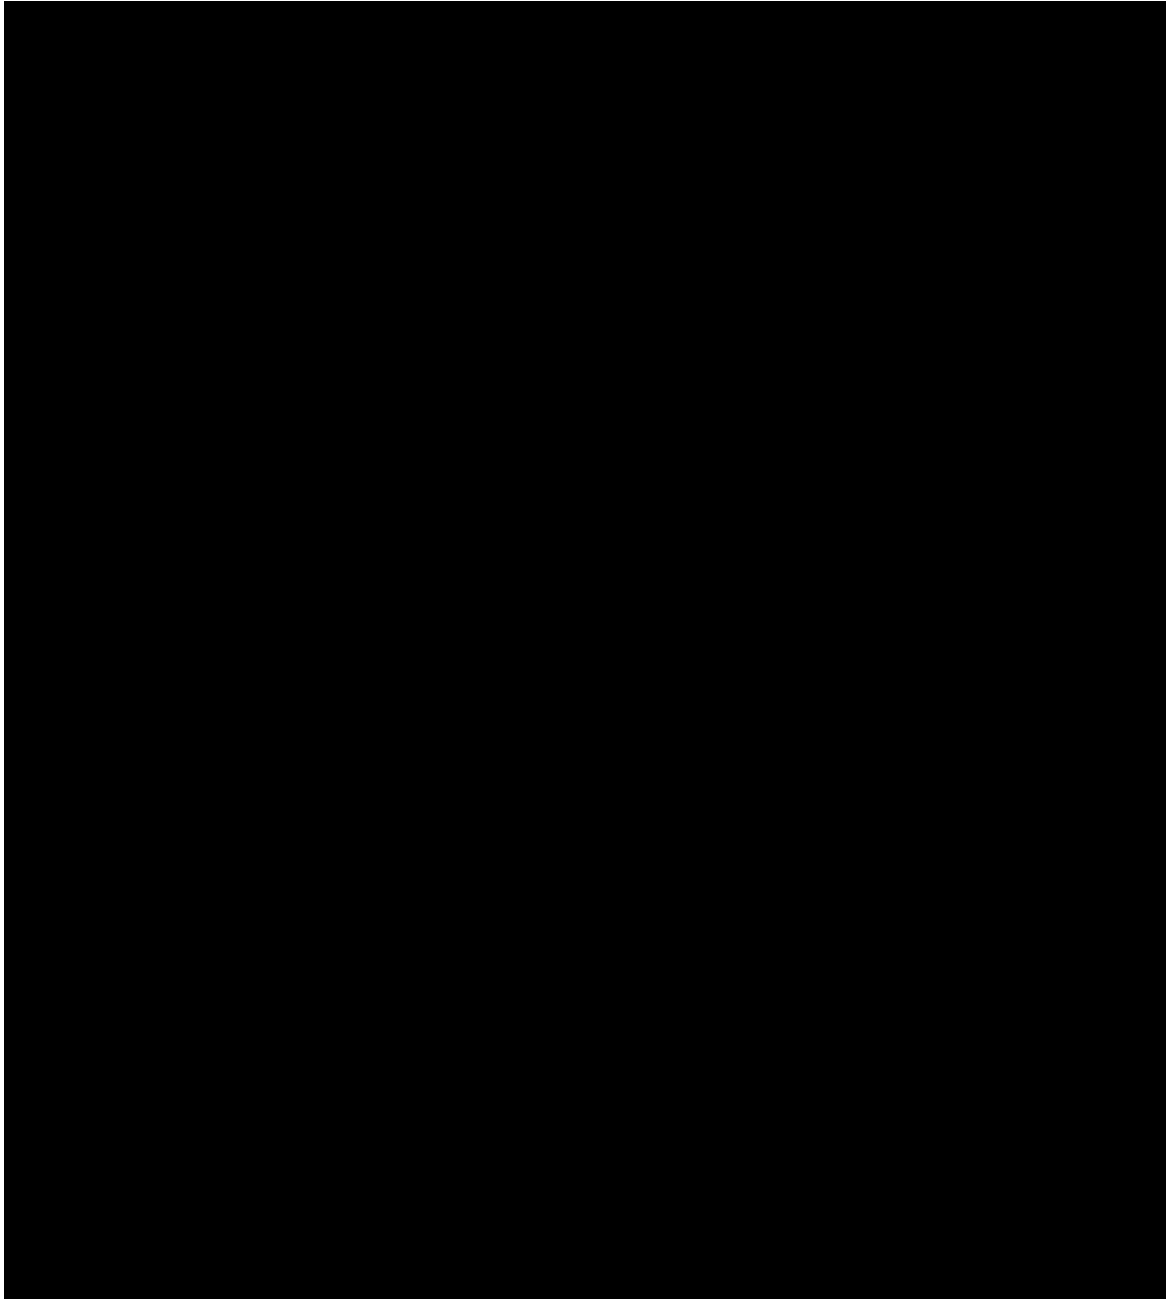
We generate 10 different *instances* using the above procedure, and the unweighted average over all products of the coefficient of variation of daily demand in units increased from ██████% for the original dataset to ██████% for the instances. Figure 7.3 clearly how this variance is strongly consolidated by the conversion to total pallets (the weighted sum of variables), with the coefficient of variation of daily demand in pallets equal to just ██████%.

We summarize the validation results for all instances and scenarios in Table G.3 on the next page. Comparing the validation results to Table 7.6 confirms that the inflated

¹We have $\text{Var}(\sum_i^n a_i X_i) = \sum_{i=1}^n a_i^2 \text{Var}(X_i) + 2 \sum_{i=1}^n \sum_{i < j \leq n} a_i a_j \text{Cov}(X_i, X_j)$.

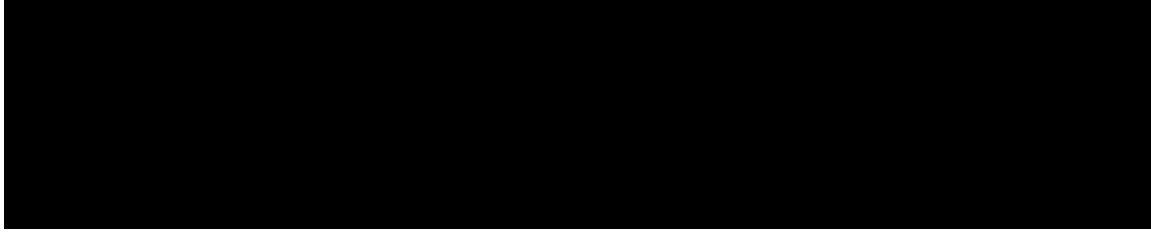
²Note that $\lfloor \dots + 0.5 \rfloor$ simply rounds to the nearest integer (for positive numbers).

Table G.3 Validation results for drops



variance is dampened by the consolidation effect. The variance of all KPIs is negligible given the amount of variance introduced.

Table G.4 Validation results for drops (average and standard deviation)



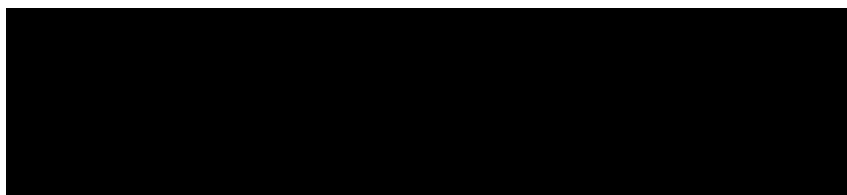
We compare the different scenarios summarized over all instances in Table G.4, since the robustness of the savings is more important than robustness of a single scenario to demand variation. The mean and standard deviation of the KPIs across instances are shown using μ and σ , respectively. Note that synchronization is the most sensitive, since filling vehicles also moderates variance. The coefficient of variation of all KPIs is lower than ■■■%, hence we conclude that our results have sufficient statistical power.

The coefficient of variation of the total daily demand in pallets is high (■■■%), hence according to Coppens [2012, p. vii] the synchronization results should indeed be robust. When the coefficient of variation is low, consistent undershoot or overshoot of the truck capacity drives transportation costs rather than consolidation effects.

Face validity

The synchronization scenario at frequency 2 most closely resembles reality. Hence, we extensively verified the results with ■■■■ (see Table G.5). All simulated KPIs were viewed as realistic by ■■■■'s project manager.

Table G.5 Face validity of simulation results



■■■■ stated that they usually use a safety stock of 6 days (whereas we use 2 days) and additionally keep safety stock for peak periods (which we were unable to forecast). Our average physical stock levels are therefore realistic. The extrapolated annual holding costs are approximately €■■■■, which was also perceived as realistic.

Bibliography

- Cachon, G. (2001, June). Managing a Retailer's Shelf Space, Inventory, and Transportation. *Manufacturing & Service Operations Management* 3(3), 211–229.
- Çetinkaya, S. and C.-Y. Lee (2000, February). Stock Replenishment and Shipment Scheduling for Vendor-Managed Inventory Systems. *Management Science* 46(2), 217–232.
- Coppens, R. (2012). *Horizontal Collaboration in Retail Supply Chains – a Study on the Benefits of Collaborative Shipping*. Thesis, Eindhoven University of Technology.
- Cruijssen, F., W. Dullaert, and H. Fleuren (2007). Horizontal Cooperation in Transport and Logistics: A Literature Review. *Transportation Journal* 46(3), 22–39.
- Daganzo, C. F. (1996). *Logistics systems analysis*. Springer.
- Hadley, G. and T. Whitin (1963). *Analysis of Inventory Systems*. Englewood Cliffs, NJ: Palgrave Macmillan.
- Hanson Logistics (2013). Reach Your Customers with the Hanson Velocities Multi-Vendor Consolidation Program.
- Higginson, J. K. (1995). Recurrent decision approaches to shipment-release timing in freight consolidation. *International Journal of Physical Distribution & Logistics Management* 25(5), 3–23.
- Johnson, D. S. (1974, June). Fast algorithms for bin packing. *Journal of Computer and System Sciences* 8(3), 272–314.
- Langevin, A., P. Mbaraga, and J. F. Campbell (1996, June). Continuous approximation models in freight distribution: An overview. *Transportation Research Part B: Methodological* 30(3), 163–188.
- Larsen, C. and G. Kiesmüller (2007, September). Developing a closed-form cost expression for an policy where the demand process is compound generalized Erlang. *Operations Research Letters* 35(5), 567–572.
- Le Blanc, H., F. Cruijssen, H. Fleuren, and M. De Koster (2006, November). Factory gate pricing: An analysis of the Dutch retail distribution. *European Journal of Operational Research* 174(3), 1950–1967.
- Lee, H. L., V. Padmanabhan, and S. Whang (1997). The Bullwhip Effect In Supply Chains. *Sloan Management Review* 38(3), 93–102.

- Mason, R., C. Lalwani, and R. Boughton (2007). Combining vertical and horizontal collaboration for transport optimisation. *Supply Chain Management: An International Journal* 12(3), 187–199.
- Mutlu, F., S. Çetinkaya, and J. H. Bookbinder (2010, February). An analytical model for computing the optimal time-and-quantity-based policy for consolidated shipments. *IIE Transactions* 42(5), 367–377.
- ORTEC (2013). ORTEC Pallet Load Building Software.
- Schneider, H. (1978, December). Methods for Determining the Re-order Point of an (s, S) Ordering Policy when a Service Level is Specified. *Journal of the Operational Research Society* 29(12), 1181–193.
- Schruben, L. W. (1980, March). Establishing the credibility of simulations. *SIMULATION* 34(3), 101–105.
- Schuijbroek, J. (2013). *Multi-vendor shipment consolidation from shared distribution centers*. Literature review, Eindhoven University of Technology.
- Shapley, L. (1953). A value for n-person games. In H. Kuhn and A. Tucker (Eds.), *Contributions to the Theory of Games II, Annals of Mathematics Studies* 28, pp. 307–317. Princeton University Press.
- Simatupang, T. M. and R. Sridharan (2002). The Collaborative Supply Chain. *The International Journal of Logistics Management* 13(1), 15–30.
- Slikker, M. (2012). Game Theory Lecture Notes. Technical report.
- Swenseth, S. R. and M. R. Godfrey (2002, May). Incorporating transportation costs into inventory replenishment decisions. *International Journal of Production Economics* 77(2), 113–130.
- Van Aken, J. E. (2005, March). Management Research as a Design Science: Articulating the Research Products of Mode 2 Knowledge Production in Management. *British Journal of Management* 16(1), 19–36.
- Van Aken, J. E., H. Berends, and H. Van der Bij (2007). *Problem Solving in Organizations: A Methodological Handbook for Business Students*. Cambridge University Press.
- Van der Vlist, P. (2007). *Synchronizing the retail supply chain*. Ph. D. thesis, Erasmus University Rotterdam.
- Van Moorsel, J. (2009). *Reduction of handling costs at the goods receiving area: The role of information integration and SSCC pallet labeling*. Thesis, Eindhoven University of Technology.