Technische Universität München
TUM School of Management

Inventory Management with Integrated Shipment Decisions and Advance Demand Information

Jana Ralfs

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung einer

Doktorin der Wirtschafts- und Sozialwissenschaften (Dr. rer. pol.)

genehmigten Dissertation.

Vorsitz:		Prof. Dr. Miriam Bird
Prüfende der Dissertation:	1.	Prof. Dr. Gudrun Kiesmüller
	2.	Prof. Dr. Robert Boute

Die Dissertation wurde am 24.05.2024 bei der Technischen Universität München eingereicht und von der TUM School of Management am 15.08.2024 angenommen.

Abstract

We consider three different inventory models in combination with shipment consolidation. In each model, stochastic demand is accumulated at a warehouse before being dispatched to customers. In two of these models, customers share advance demand information, which is taken into account at the time of shipment. We jointly optimize the inventory and shipment policy parameters and also provide structural insights into the optimal outbound shipment quantities for a time-based dispatch scheme.

Firstly, we present a model for a one-warehouse-multi-retailer inventory system that receives stochastic demand. The warehouse satisfies retailer demand using a hybrid time-and-quantity-based shipment consolidation policy. Shipments are dispatched according to a time-based schedule; however, additional quantity-based shipments may occasionally be dispatched in between if the number of accumulated demands reaches a specific consolidation quantity to reduce retailer waiting time. We derive the probability mass function of the inventory level at each retailer, enabling efficient computation of the system's inventory and shipment costs. After evaluating various inventory and shipment policies, we demonstrate how the optimal policy parameters can be computed. A numerical study shows that using the pure time-based or pure quantity-based policy instead of the dominating hybrid policy can be implemented without significant total cost increases in most instances.

Secondly, an inventory model is introduced that incorporates shipment decisions and advance demand information to analyze a flexible time-based shipment policy. The orders of a production facility are accumulated at the warehouse, and shipments leave based on a periodic dispatch scheme. Even though orders arrive with a due date, the warehouse is allowed to fulfill these orders ahead of this due date to increase the utilization of transportation means. We provide analytical, approximate expressions for the inventory and shipment costs, enabling the evaluation of various inventory and shipment policies. Moreover, we demonstrate how to optimize the policy parameters. Our approximation reaches sufficient accuracy and is validated through comparison with results obtained by simulation. Our findings indicate that advance demand information not only reduces the safety stock needed at the warehouse but also enables cost savings by integrating flexible deliveries into already familiar and adopted shipment policies.

Finally, we formulate a Markov decision process to model an inventory system receiving advance demand information and where shipments are dispatched according to a time-based schedule. The aim is to answer the question about the optimal outbound shipment quantities when flexible deliveries are allowed. Since large-scale instances cannot be solved to optimality as the number of dimensions increases, a deep reinforcement learning algorithm is implemented to observe the structure of the (near-)optimal shipment policy. Building on this insight, we propose an approximated threshold policy, which is compared to simple benchmark policies and shows great performance across a wide range of instances.

Acknowledgements

A formative chapter is coming to an end. As I reflect upon the journey that has led me to this moment, I am aware that this path was not traveled alone. With the immense support and encouragement of various people, I would have the honor and privilege of composing this thesis.

First and foremost, I would like to thank my supervisor Professor Gudrun Kiesmüller. I am thankful for the continuous support and guidance she provided to me throughout my academic and professional journey. Under her mentorship, I had the privilege of expanding my skills in scientific research, with a particular focus on operations management. The memories of the various business trips we experienced together remain an appreciated part of my journey. Especially, I will always value the events organized by the International Society on Inventory Research. I am immensely grateful to Gudrun Kiesmüller for the great years we spent together in Magdeburg and Heilbronn.

I would like to extend my sincere gratitude to Filip and Professor Johan Marklund for our successful collaboration and their warm hospitality during my stay in Lund. Furthermore, I would like to thank Dai for the fruitful collaboration. Additionally, my sincere thanks go to Professor Robert Boute for participating in the examination committee and to Professor Miriam Bird for chairing it.

I would also like to thank my colleagues at the Chair of Operations Management at the Otto-von-Guericke University Magdeburg and the Center of Digital Transformation at the Technical University of Munich. I would like to give special recognition to Florian, Julia, and Ung for making my time truly exceptional.

A heartfelt thanks go to my family. I am profoundly grateful for your continuous emotional support.

Contents

List of Figures

List of Tables

List of Abbreviations and Symbols

Abbreviations

3PL	Third-party logistics service provider
ADI	Advance demand information
ATH	Approximate threshold
DRL	Deep reinforcement learning
FCFS	First-come-first-serve
OWMR	One-warehouse-multiple-retailer
PASTA	Poisson arrivals see time averages
PMF	Probability mass function
ReLU	Rectified linear unit
ShipOp1	Shipment Option 1
ShipOp2	Shipment Option 2
VI	Value Iteration

Symbols

b_0	Realization of $B_0(\tau' - T_I)$
b_0^m	Realization of $B_0^m(\tau' - T_I)$, $m \in \{1, 2, \ldots, \mathcal{M}\}$

$B(m_1, m_2)$	Beta function with parameters m_1 and m_2
$B_0(t)$	Backorders at the warehouse at time t
$B_0^m(t)$	Backorders at the warehouse destined for retailer group m at time t, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$Beta(p, q)$	Beta distribution with parameters p and q
$Betabin(k, \alpha, \beta)$	Beta-binomial distribution with parameters k, α and β
$Bin(n, p)$	Binomial distribution with parameters n and p
c_1	Variable shipment costs per unit reserved for the primary transportation option
c_2	Variable shipment costs per unit shipped by the alternative transportation option
c_q^m	Variable shipment costs per unit shipped in a quantity-based shipment to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
c_t^m	Variable shipment costs per unit shipped in a time-based shipment to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$c(\bar{S}_t, x_t)$	Expected total costs per period (depending on state $\bar{S}_t$ and decision x_t)
$c^{det}(\bar{S}_t, x_t)$	Deterministic part of the total cost per period (depending on state $\bar{S}_t$ and decision x_t)
$c^{rand}(\bar{S}_t, x_t)$	Random part of the total cost per period (depending on state $\bar{S}_t$ and decision x_t)
$c_t^{early}(x_t)$	Early-delivery costs per shipment period t (depending on decision x_t), $t \in \mathcal{T}$
$c_t^{late}(x_t)$	Late-delivery costs per shipment period t (depending on decision x_t), $t \in \mathcal{T}$
$c_t^{shipment}(x_t)$	Additional shipment costs for the alternative transportation option per shipment period t (depending on decision x_t), $t \in \mathcal{T}$
$C(\Omega(S), V, M(t_n))$	Inventory costs for Case $S > 0$ (depending on the random variables $\Omega(S)$, V and $M(t_n)$)

$\tilde{C}(\Omega(S), V)$	Inventory costs for Case $S \leq 0$ (depending on the random variables $\Omega(S)$ and V)
$C_i(\Omega(S), V)$	Inventory costs for Case $S > 0$ and Situation i (depending on the random variables $\Omega(S)$ and V), $i \in \{A, B, C, D, E, F, G\}$
$C_i(\Omega(S), V, M(t_n))$	Inventory costs for Case $S > 0$ and Situation i (depending on the random variables $\Omega(S)$, V and $M(t_n)$), $i \in \{A, B, C, D, E, F, G\}$
$\bar{C}_i(\Omega(S))$	Inventory costs for Situation i without flexible deliveries (depending on the random variable $\Omega(S)$), $i \in \{A, G\}$
Cap	Reserved transportation capacity
Cap^*	Optimal reserved transportation capacity
Cap^l	Lower bound on Cap^*_{app}
Cap^u	Upper bound on Cap^*_{app}
d_{max}	Maximum value of the random uniformly distributed demand
d_{min}	Minimum value of the random uniformly distributed demand
d_t^n	Realization of D_t^n, $n \in \{0, 1, \ldots, T-1\}$
$D(s,t)$	Orders at the warehouse during the time interval $(s,t]$ in the single-echelon inventory system, $s < t$
$D_0(s,t)$	Demand at the warehouse during the time interval $(s,t]$ where the demand at time t is included, $s < t$
$\tilde{D}_0(s,t)$	Demand at the warehouse during the time interval $[s,t)$ where the demand at time s is included, $s < t$
$D^m(s,t)$	Demand at retailer group m during the time interval $(s,t]$, $s < t$, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$D^{\overline{m}}(s,t)$	Demand at all retailer groups excluding group m during the time interval $(s,t]$, $s < t$, $m \in \{1, 2, \ldots, \mathcal{M}\}$
D_t^n	Random number of ordered units at the end of period $t+n$, which are due in period $t+n+L_d+1$, $t \in \mathcal{T}$, $n \in \{0, 1, \ldots, T-1\}$
e	Early-delivery costs per unit and time unit/period at the warehouse

e	Euler's number, mathematical constant
$f(z)$	Probability density function of $M(t_n)$ when $S > 0$
$f(x, y, z)$	Joint probability density function of $\Omega(S)$, V and $M(t_n)$ when $S > 0$
$\tilde{f}(x, y)$	Joint probability density function of $\Omega(S)$ and V when $S \leq 0$
$f_i(x, y)$	Joint probability density function of $\Omega(S)$ and V for Situation i when $S > 0$, $i \in \{A, B, C, D, E, F, G\}$
$f_i(x, y, z)$	Joint probability density function of $\Omega(S)$, V and $M(t_n)$ for Situation i when $S > 0$, $i \in \{A, B, C, D, E, F, G\}$
$\tilde{F}(x, y)$	Joint cumulative distribution function when $S \leq 0$
$F^{\lambda}(k)$	Cumulative distribution function of the Poisson distribution with parameter λ
$F_i(x, y)$	Joint cumulative distribution function of Situation i when $S > 0$, $i \in \{A, B, C, D, E, F, G\}$
$g^x(t)$	Density of an Erlang distribution with parameters (x, λ)
$G^x(t)$	Cumulative distribution function of an Erlang distribution with parameters (x, λ)
h	Holding costs rate per unit and per time unit at the warehouse in the single-echelon inventory system
h_i	Holding costs rate per unit and per time unit at stock point i in the multi-echelon inventory system, $i \in \{0, 1, \ldots, \mathcal{N}\}$
$Hyp(N, K, n)$	Hypergeometric distribution with parameters N, K and n
ip	Realization of $IP_0(\tau' - L_0 - T_I)$
I	Target network update interval
I^m	Number of units in a time-based shipment to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$IL(t)$	Inventory level at the warehouse at time t in the single-echelon inventory system

$IL^-(t)$	Backorders at the warehouse at time t in the single-echelon inventory system
IL_0	Inventory level of the unreserved stock at the warehouse in the multi-echelon inventory system
IL_0^+	Unreserved stock on hand at the warehouse in the multi-echelon inventory system
IL_0^-	Backorders at the warehouse in the multi-echelon inventory system
IL_i	Inventory level at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$
IL_i^+	Stock on hand at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$
IL_i^-	Backorders at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$
$IP(t)$	Inventory position at the warehouse at time t in the single-echelon inventory system
$IP_0(t)$	Inventory position at the warehouse at time t in the multi-echelon inventory system
J^m	Number of quantity-based shipments within a shipment interval T^m to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
K	Number of units that qualify for shipment to all retailer groups between two arbitrary time-based shipments
$K(t)$	Remaining units that cannot be shipped at time t due to a lack of stock at the warehouse or limited transportation capacity
K^m	Number of units that qualify for shipment to retailer group m between two arbitrary time-based shipments, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$K^0(t)$	$\lim\limits_{Cap \to 0} K(t)$
$K^\infty(t)$	$\lim\limits_{Cap \to \infty} K(t)$
l	Late-delivery costs per unit and time unit/period at the warehouse

$\mathcal{L}$	Loss function
$\boldsymbol{L}$	Vector representation of L_i, $\boldsymbol{L} = (L_1, L_2, \dots, L_{\mathcal{N}})$
L_0	Constant supply lead time from an outside supplier to the warehouse in the multi-echelon inventory system
L_i	Constant transportation time from the warehouse to retailer i in the multi-echelon inventory system, $i \in \{1, 2, \dots, \mathcal{N}\}$
L_d	Constant demand lead time in the single-echelon inventory system
L_s	Constant supply lead time from an outside supplier to the warehouse in the single-echelon inventory system
N	Total number of demands and orders
$\mathcal{N}$	Number of non-identical retailers
$\mathbb{N}$	Natural numbers, $\mathbb{N} = \{1, 2, 3, \dots\}$
$\mathbb{N}_0$	Natural numbers including zeros, $\mathbb{N}_0 = \{0, 1, 2, \dots\}$
m	Realization of M
M	Shipment quantity from the warehouse to the production facilities, $M \equiv M(t)$
$M(t)$	Shipment quantity from the warehouse to the production facility at time t
$\mathcal{M}$	Number of retailer groups
n	Priority of the considered n^{th} unit
O_i	Number of units outstanding for retailer i, $i \in \{1, 2, \dots, \mathcal{N}\}$
$\tilde{O}_i$	Number of units outstanding for retailer i ordered after $\tau - L_i$, $i \in \{1, 2, \dots, \mathcal{N}\}$
$\widehat{O}_i$	Number of units outstanding for retailer i ordered before $\tau - L_i$, $i \in \{1, 2, \dots, \mathcal{N}\}$
$\widehat{O}^m$	Number of units outstanding for retailers in retailer group m ordered before $\tau - L_i$, $m \in \{1, 2, \dots, \mathcal{M}\}$, $i \in \{1, 2, \dots, \mathcal{N}\}$

$\boldsymbol{p}$	Vector representation of p_i, $\boldsymbol{p} = (p_1, p_2, \ldots, p_{\mathcal{N}})$
p_i	Backorder costs rate per unit and per time unit at retailer i, $i \in \{1, 2, \ldots, \mathcal{N}\}$
$\mathcal{P}_i(x)$	Probability mass function of X conditioned on Case i, $Pr(X = x \vert \text{Case } i)$, $i \in \{\text{A,B}\}$
$\mathcal{P}_i(x\vert y)$	Probability mass function of X conditioned on $Y = y$ and Case i, $Pr(X = x\vert Y = y, \text{Case } i)$, $i \in \{\text{A,B}\}$
$\mathcal{P}_i(x, y)$	Joint probability mass function of X and Y conditioned on Case i, $Pr(X = x, Y = y\vert \text{Case } i)$, $i \in \{\text{A,B}\}$
$Pr(X = x)$	Probability mass function of X
$Poisson(\lambda)$	Poisson distribution with parameter λ
q	Priority of the last unit in a quantity-based shipment including the considered n^{th} unit
Q	Order quantity at the warehouse in the single-echelon inventory system
$\mathbf{Q}$	Vector representation of Q^m, $\mathbf{Q} = (Q^1, Q^2, \ldots, Q^{\mathcal{M}})$
Q_0	Order quantity at the warehouse in the multi-echelon inventory system
Q^m	Consolidation quantity to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$
Q^{m*}	Optimal consolidation quantity to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\mathcal{Q}^m$	Set of feasible consolidation quantities to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\mathbf{Q}^*$	Vector representation of Q^{m*}, $\mathbf{Q}^* = (Q^{1*}, Q^{2*}, \ldots, Q^{\mathcal{M}*})$
$\mathbb{Q}_{>0}$	Positive rational numbers, all positive fractions whose numerators and denominators consist of integers
r_c	Shipment costs ratio

r_{Cap}	Capacity ratio
R	Reorder level at the warehouse in the single-echelon inventory system
R_0	Reorder level at the warehouse in the multi-echelon inventory system
R^*	Optimal reorder level at the warehouse obtained by simulation in the single-echelon inventory system
R^l	Lower bound on R^*_{app}
R^u	Upper bound on R^*_{app}
R^*_0	Optimal reorder level at the warehouse in the multi-echelon inventory system
R^l_0	Lower bound on R^*_0
R^u_0	Upper bound on R^*_0
R^*_{app}	Optimal reorder level at the warehouse obtained by the approximation in the single-echelon inventory system
R^*_{asm}	Optimal reorder level at the warehouse without the policy assumption obtained by simulation in the single-echelon inventory system
S	Base-stock level at the warehouse in the single-echelon inventory system
$\mathbf{S}$	Vector representation of S_i, $\mathbf{S} = (S_1, S_2, \ldots, S_{\mathcal{N}})$
S_i	Base-stock level at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$
$\bar{S}_t$	Pre-decision state, state at the beginning of shipment period t, $\bar{S}_t = (w_t, y_t^1, \ldots, y_t^{L_d})$, $t \in \mathcal{T}$
$\mathcal{S}_t$	Set of all feasible pre-decision states in shipment period t, $t \in \mathcal{T}$
$\mathbf{S}^*$	Vector representation of S^*_i, $\mathbf{S}^* = (S^*_1, S^*_2, \ldots, S^*_{\mathcal{N}})$

S_i^*	Optimal base-stock level at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$
$\bar{S}_t^x$	Post-decision state, state after making decision x_t in shipment period t, $\bar{S}_t^x = (\tilde{w}_t, \tilde{y}_t^1, \ldots, \tilde{y}_t^{L_d})$, $t \in \mathcal{T}$
t_a	Time when the replenishment order including the considered unit arrives at the warehouse, $t_a = t_r + L_s$
t_d	Time when the demand for the considered unit occurs at the warehouse, $t_o = t_r + \Omega(S) + L_d$
t_n	n^{th} shipment day, $t_n = nT$
t_o	Time when the order for the considered unit occurs at the warehouse, $t_o = t_r + \Omega(S)$
t_r	Time when the warehouse just placed a replenishment order including the considered unit
t_{ro}	Time when the retailer orders the last unit with priority q in a quantity-based shipment including the considered n^{th} unit
t_{wo}	Time when the warehouse orders the last unit with priority q in a quantity-based shipment including the considered n^{th} unit
T	Shipment interval in the single-echelon inventory system
$\mathbf{T}$	Vector representation of T^m, $\mathbf{T} = (T^1, T^2, \ldots, T^{\mathcal{M}})$
$\mathcal{T}$	Set of decision periods in the single-echelon inventory system, $\mathcal{T} = \{0, 1T, 2T, \ldots\}$
T_I	Uniform random variable on the interval $[0, T^m]$, $m \in \{1, 2, \ldots, \mathcal{M}\}$
T^*	Optimal shipment interval obtained by simulation in the single-echelon inventory system
T^l	Lower bound on T^*_{app}
T^u	Upper bound on T^*_{app}
T^m	Shipment interval to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$

T^{m*}	Optimal shipment interval to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\mathcal{T}^m$	Set of feasible shipment intervals to retailer group m in the multi-echelon inventory system, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\mathbf{T}^*$	Vector representation of T^{m*}, $\mathbf{T}^* = (T^{1*}, T^{2*}, \ldots, T^{\mathcal{M}*})$
T^*_{app}	Optimal shipment interval obtained by the approximation in the single-echelon inventory system
T^*_{asm}	Optimal shipment interval without the policy assumption obtained by simulation in the single-echelon inventory system
$TC(R, T, Cap)$	Expected total cost per time unit in the single-echelon inventory system (depending on the decision variables R, T and Cap)
$TC_{sim}(R, T)$	Expected total cost per time unit obtained by simulation in the single-echelon inventory system (depending on the decision variables R and T)
$TC_M(R_0, \mathbf{S}, \mathbf{T}, \mathbf{Q})$	Expected total cost per time unit in the multi-echelon inventory system (depending on the decision variables R_0, $\mathbf{S}$, $\mathbf{T}$ and $\mathbf{Q}$)
$TIC(R, T, Cap)$	Expected inventory costs per time unit in the single-echelon inventory system (depending on the decision variables R, T and Cap)
$TSC(R, T, Cap)$	Expected shipment costs per time unit in the single-echelon inventory system (depending on the decision variables R, T and Cap)
$u(t)$	Density function of a uniform distributed random variable on the interval $[0, T]$
$U(t)$	Cumulative distribution function of a uniform distributed random variable on the interval $[0, T]$
$Uniform(a, b)$	Uniform distribution with parameters a and b
V	Shipment delay due to the shipment policy
$V(\bar{S}_t)$	Value function of the pre-decision state $\bar{S}_t$

$V_x(\bar{S}_t^x)$	Value function of the post-decision state $\bar{S}_t^x$
$V_x^\theta(\bar{S}_t^x)$	Value function of the post-decision state $\bar{S}_t^x$ obtained by the deep neural network θ
w_t	Demands at the beginning of shipment period t which are already due, $t \in \mathcal{T}$
$\tilde{w}_t$	Demands after making decision x_t in shipment period t which are already due, $t \in \mathcal{T}$
$\bar{w}_t$	Remaining demands at the beginning of shipment period t which cannot be shipped by the reserved capacity Cap and which are already due, $t \in \mathcal{T}$
W_i	Reserved stock on hand at the warehouse destined for retailer, $i \in \{1, 2, \ldots, \mathcal{N}\}$
x_t	Shipment quantity at the beginning of shipment period t, $t \in \mathcal{T}$
$X_{\beta_j^i}$	Number of demands among β_j^i demands that occur before the last unit with priority q was ordered in Interval j for Case i, $i \in \{\mathsf{A}, \mathsf{B}\}$, $j \in \{1, 2, 3\}$
$X_{\beta_4^i}$	Number of demands among β_4^i demands that occur after the last unit with priority q was ordered but before $\tau' - T_I - L_0$ for Case i, $i \in \{\mathsf{A}, \mathsf{B}\}$
$\mathcal{X}_t$	Set of potential decisions at the beginning of shipment period t, $\mathcal{X}_t = \{0, 1, \ldots, w_t + \sum_{i=1}^{L_d} y_t^i\}$, $t \in \mathcal{T}$
y_t^n	Orders at the beginning of shipment period t which are due in period $t+n$, $t \in \mathcal{T}$, $n \in \{1, 2, \ldots, L_d\}$
$\tilde{y}_t^n$	Orders after making decision x_t in shipment period t which are due in period $t+n$, $t \in \mathcal{T}$, $n \in \{1, 2, \ldots, L_d\}$
$\bar{y}_t^n$	Remaining orders in the beginning of shipment period t which are due in period $t+n$ and cannot be shipped by the reserved capacity Cap, $t \in \mathcal{T}$, $n \in \{1, 2, \ldots, L_d\}$
$\mathbb{Z}$	Integers, $\mathbb{Z} = \{\ldots, -3, -2, -1, 0, 1, 2, 3, \ldots\}$

$\boldsymbol{\alpha^i}$	Vector representation of α_j^i for Case i, $\boldsymbol{\alpha^i} = (\alpha_1^i, \alpha_2^i, \alpha_3^i)$, $i \in$ {A,B}
α_j^i	Demand realization of retailer group m during Interval j for Case i, $i \in$ {A,B}, $j \in \{1, 2, 3\}$
$\boldsymbol{\beta^i}$	Vector representation of β_j^i for Case i, $\boldsymbol{\beta^i} = (\beta_1^i, \beta_2^i, \beta_3^i)$, $i \in$ {A,B}
β_j^i	Demand realization of the all retailer groups excluding retailer group m during Interval j for Case i, $i \in$ {A,B}, $j \in \{1, 2, 3\}$
γ	Discount factor
$\boldsymbol{\gamma^i}$	Vector representation of γ_j^i for Case i, $\boldsymbol{\gamma^i} = (\gamma_1^i, \gamma_2^i, \gamma_3^i)$, $i \in$ {A,B}
γ_j^i	System demand realization during Interval j for Case i, $\gamma_j^i = \alpha_j^i + \beta_j^i$, $i \in$ {A,B}, $j \in \{1, 2, 3\}$
$\delta(\psi \vert ip)$	Number of system demands that occurs during the interval $(\tau' - T_I - L_0, t_{wo}]$ if $\tau' - T_I - L_0 < t_{wo}$, or number of system demands that occur during the interval $[t_{wo}, \tau' - T_I, L_0)$ if $\tau' - T_I - L_0 > t_{wo}$ (depending on the realizations of ψ and ip)
ϵ	Probability of choosing a random, feasible decision
ε	Small positive number close to zero
θ	Deep neural network
θ^{best}	Best-found deep neural network parameters
θ^{target}	Target deep neural network
Θ_i	Set of feasible combinations of values in the support of the joint probability mass function for Case i, $i \in$ {A,B}
λ	Expected demand per time unit at the warehouse in the single-echelon inventory system
$\boldsymbol{\lambda}$	Vector representation of λ_i in the multi-echelon inventory system, $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_{\mathcal{N}})$
λ_i	Expected demand per time unit at retailer i in the multi-echelon inventory system, $i \in \{1, 2, \ldots, \mathcal{N}\}$

λ^m	Expected demand per time unit at retailer group m in the multi-echelon inventory system, $\lambda^m = \sum_{i \in \Omega^m} \lambda_i$, $m \in \{1, 2, \dots, \mathcal{M}\}$
Λ	Expected total demand per time unit at the warehouse in the multi-echelon inventory system, $\Lambda = \sum_{i=1}^{\mathcal{N}} \lambda_i$
μ_j^i	Threshold (i, j) of the observed threshold policy, $i \in \{1, 2, \dots, L_d\}$, $j \in \{0, 1, \dots, L_d\}$
μ_j^{ATH}	Threshold j of the approximate threshold policy, $j \in \{0, 1, 2\}$
π	Policy, function that maps a pre-decision state $\bar{S}_t$ to a decision x_t
π^*	Optimal policy that minimizes the expected total discounted period costs C_t
π^{ATH}	Approximate threshold poilicy
π^{DRL}	Deep reinforcement learning policy
π^{GRDY}	Greedy policy
π^{LAZY}	Lazy policy
π^{UPTO}	Dispatch-up-to policy
$\pi_{L_d}^{TH}$	Observed threshold policy where the number of thresholds depend on L_d
Π	Set of feasible policies π
τ	Arbitrary time point
τ'	Latest time an order can be placed by retailer i so that this order is available at retailer i by τ, $\tau - L_i$, $i \in \{1, 2, \dots, \mathcal{N}\}$
$\phi(\psi \vert ip)$	Number of system demands that occurs during the interval $(t_{wo}, t_{ro}]$ (depending on the realizations ψ and ip)
ϕ_t	Realization of vector Φ_t, $\phi_t = (d_t^0, d_t^1, \dots, d_t^{T-1})$, $t \in \mathcal{T}$
Φ_t	Vector representation of D_t^n, $\Phi_t = (D_t^0, D_t^1, \dots, D_t^{T-1})$, $t \in \mathcal{T}$
ψ	Realization of $\Psi(q)$

$\Psi(q)$	Number of system demands that occurs during the interval $(\tau' - T_I - L_0, t_{ro}]$ (depending on q)
$\boldsymbol{\omega}_q$	Vector representation of ω_q^m, $\boldsymbol{\omega}_q = (\omega_q^1, \omega_q^2, \ldots, \omega_q^{\mathcal{M}})$
ω_q^m	Fixed shipment costs per quantity-based shipment to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\boldsymbol{\omega}_t$	Vector representation of ω_t^m, $\boldsymbol{\omega}_t = (\omega_t^1, \omega_t^2, \ldots, \omega_t^{\mathcal{M}})$
ω_t^m	Fixed shipment costs per time-based shipment to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
$\omega(Cap)$	Fixed shipment costs per shipment for time-based shipments (depending on the decision variable Cap)
$\Omega(x)$	Length of the time interval between the replenishment moment at the warehouse and the moment when the x^{th} order from the production facility arrives at the warehouse
Ω^m	Set of retailers belonging to retailer group m, $m \in \{1, 2, \ldots, \mathcal{M}\}$
x^+	$\max(0, x)$
x^-	$\max(0, -x)$
$\mathrm{mod}_n(x)$	$x - n\lfloor \frac{x}{n} \rfloor$
$\underset{n,m}{\mathrm{mod}}(x)$	$x + km$, where k is an integer such that $n < x + km \leq n + m$
$\lfloor x \rfloor$	$\max(k \in \mathbb{Z} \| k \leq x)$
$\lceil x \rceil$	$\min(k \in \mathbb{Z} \| k \geq x)$
$\lfloor x \rceil$	$\lfloor x + 0.5 \rfloor$

Chapter

1 Introduction

As supply chains grow in complexity, it becomes increasingly important to hedge against sources of uncertainties, e.g., demand, supply, and lead time uncertainties. Keeping safety stock is one strategy to protect against these uncertainties and prevent stockouts that could result in customer dissatisfaction or production losses. However, keeping inventory on stock causes significant holding costs that can reduce profitability if not managed effectively. Therefore, deciding on the number of units to stock is a critical decision for any business, as it directly impacts both operational efficiency and financial health. Therefore, finding the right balance between holding sufficient inventory to meet demand and minimizing inventory holding costs is far from a trivial task. Consequently, inventory management has been a pivotal area of practice and research for decades.

Moreover, environmental awareness has increased significantly in recent years, underlining the growing importance of designing sustainable supply chains and distribution systems. OECD (2023) emphasizes this objective by highlighting the challenging trade-off arising from the simultaneous increase in transport demand and the need to reduce CO_2 emissions. This is particularly relevant considering that the transport sector is responsible for 23 % of global energy-related CO_2 emissions (OECD, 2023). Additionally, Doherty and Hoyle (2009) point out that the average load factor of the maximum gross weight is only 57 %, without taking into account empty truck runs. To achieve a potential reduction of up to 72 % in freight emissions compared to 2015, OECD (2021) suggests transformations of current processes, such as enhancing freight consolidation and strengthening collaborations within supply chains, leading, among other things, to higher load factors.

However, in research, inventory and shipment decisions are often analyzed separately due to reasons of complexity. Consequently, when evaluating inventory policies, it is common to assume that demand is satisfied immediately. Although replenishment policies can be strategically applied to aggregate demand from downstream stock points into batches, these demand batches are assumed to be satisfied directly, and in cases of stockouts, partial deliveries are made. Any unfulfilled demand is satisfied as soon as inventory is replenished at the upstream stock point, see, e.g., Axsäter (2000). In the area of transportation science, assumptions are often reversed. For instance, the stochastic inventory problem is characterized by the fact that demand from several downstream stock points is satisfied within a single route, and the inventory levels at these downstream stock points are taken into account. However, it is assumed that there is ample supply at the upstream stock point from where the routes are planned, ensuring that all requirements can be met. As a result, stockouts on shipment days are not considered within the analysis, as exemplified by Sonntag et al. (2023).

Hence, there has been a growing body of research in the field of inventory management with integrated shipment decisions for more than 30 years. Nevertheless, several questions within this field remain unanswered. Furthermore, Malmberg and Marklund (2023) recently emphasize that the joint analysis of inventory and shipment decisions in distribution systems is significant for sustainable and economical actions.

Finally, enhancing collaborations within the supply chains, such as exchanging information, presents a promising strategy to further improve shipment and inventory decisions. For instance, information about future demand can be shared, leading to increased utilization of transportation means and reducing demand uncertainties. Advance demand information (ADI) can be efficiently modeled by a demand lead time as demonstrated in Wang and Toktay (2008).

Based on the aforementioned facts, we contribute to the field of inventory control with integrated shipment decisions and ADI by deriving new mathematical models to analyze systems with different shipment consolidation strategies. On the one hand, we analyze continuous review inventory systems integrating heuristic shipment consolidation strategies. For these systems, the aim is to find optimal inventory and shipment parameters to minimize the total cost of the system. On the other hand, we derive a near-optimal outbound shipment strategy for a periodic inventory system, considering a positive demand lead time, a time-based dispatch strategy, and ample supply at the distribution center.

Before outlining the contribution and methodologies, and stating the research questions answered in this work, we establish essential terminology and fundamentals required for modeling various inventory systems.

1.1 Terminology and Fundamentals of Inventory Management

The objective of inventory management is to develop control strategies that minimize the costs within the inventory system while ensuring the availability of sufficient stock on hand to receive a high customer service level. To balance both objectives, the flow of goods within the inventory system has to be managed. That encompasses not only the questions of when and how much to order, but it is equally important to consider the questions of when and how much to dispatch.

The purpose of this chapter is to offer a concise introduction to inventory control theory, in particular, to introduce the fundamental definitions and concepts that are relevant to this thesis. We discuss **structures** of inventory systems, the definition of **supply lead times**, types of **review periods** and **customer demand**. Additionally, we present established **replenishment, dispatch and allocation policies** within inventory systems. We also provide a short summary of **demand lead times** in inventory systems. Lastly, we introduce measures used to evaluate the **performance** of inventory systems. The contents presented in this chapter are based on the textbooks Tempelmeier (2006) and Axsäter (2015). For a broader overview of the field of inventory management, we refer to Tempelmeier (2006), Axsäter (2015) and Silver et al. (2016).

1.1.1 Structures of Inventory Systems

To manage the flow of goods through an inventory system, it is substantial to define the topology, which encompasses the number of stock points linked together and the flow of goods and information through this system. In Figure 1.1, we present an inventory system consisting of a single stock point, commonly refereed to as a single-echelon inventory system. The stock point is illustrated as a triangle, whereas material and information flow are represented as solid arrows and dashed arrows, respectively. The stock point receives customer demand and places orders at an outside supplier, represented by dashed arrows. Solid arrows symbolize replenishment deliveries and the dispatch of units to the customers. It is noteworthy that the outside supplier and the customers are not explicitly modeled as stock points.

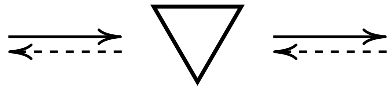

Figure 1.1: A single-echelon inventory system

Inventory systems encompassing several connected stock points at different stages are referred to as multi-echelon inventory systems. These multi-echelon inventory systems can be categorized based on the links between the stock points, in particular into serial, convergent, divergent, or general inventory systems.

In a serial inventory system, each stock point has at most one immediate predecessor and at most one immediate successor. The most upstream stock point is characterized by ample supply, whereas the most downstream stock point is of substantial importance since it receives random demand from external customers. For example, a serial inventory system may consist of one central warehouse, one local warehouse, and one retailer. Figure 1.2 provides an illustration of a three-echelon serial inventory system.

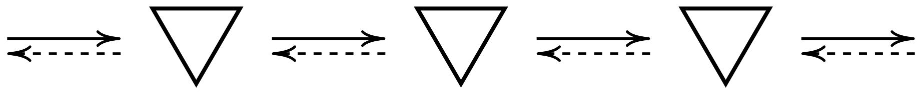

Figure 1.2: A serial multi-echelon inventory system

As inventory systems grow in complexity, multiple predecessors or successors may occur. In the case of a pure convergent inventory system, each stock point is linked to multiple immediate predecessors but, at most, to one immediate successor. These systems are used to model assembly and production environments, where various components are needed to create a final product. One example where a product is assembled out of three distinct components is illustrated in Figure 1.3.

Conversely, in a pure divergent inventory system each stock point is connected to, at most, one predecessor but may have multiple successors. One specific structure discussed in the literature is the one-warehouse-multiple-retailer (OWMR) distribution inventory system. In this system, several retailers (or local warehouses) fulfill random customer demand and place orders at a central warehouse. The central warehouse, in turn, satisfies

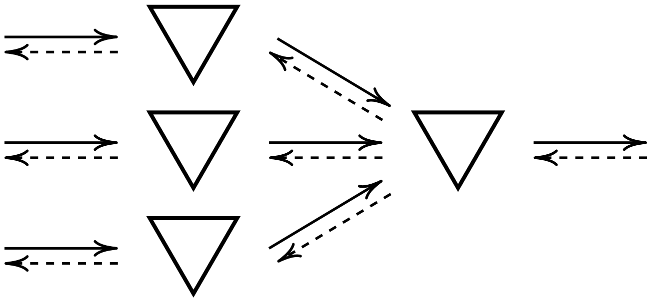

Figure 1.3: A convergent multi-echelon inventory system

retailers' orders and replenishes its own inventory from an external supplier with ample stock. Such an inventory system benefits from the pooling effect at the central warehouse because it allows for a reduction in inventories without compromising the performance of the system. Figure 1.4 represents an OWMR distribution system with three retailers.

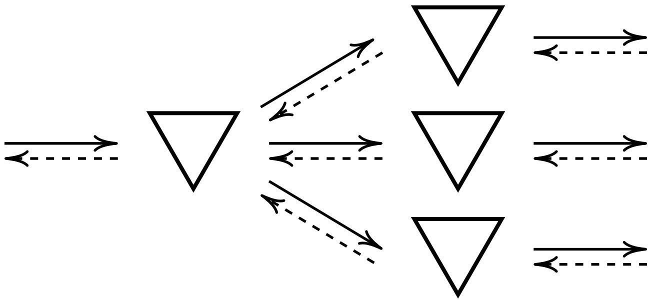

Figure 1.4: A divergent multi-echelon inventory system

In real-world scenarios, the topology of supply chains or production systems is often significantly more complex, and each stock point may be linked to multiple successors and predecessors. An example of a general multi-echelon inventory system is illustrated in Figure 1.5.

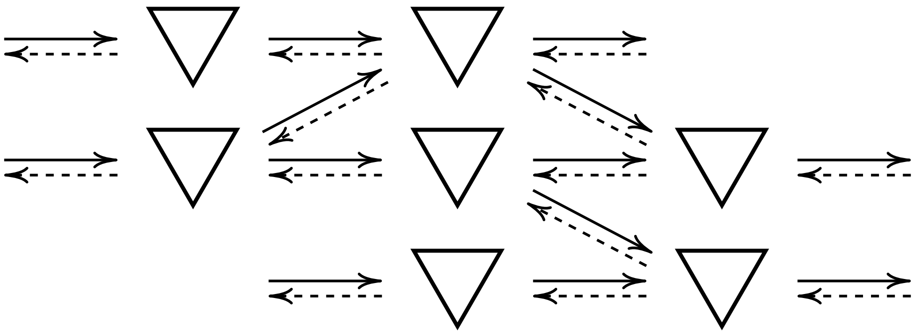

Figure 1.5: A general multi-echelon inventory system

A mathematical analysis may often not be achievable when inventory problems become highly complex. Therefore, it is essential to capture the key characteristics and trade-offs of a complex inventory problem by incorporating them into tractable models. For this purpose, depending on the aspects to be examined, single-echelon and OWMR inventory systems are often considered in the literature. For instance, single-echelon inventory systems may be used to obtain the optimal replenishment structure for an inventory system with ADI and flexible deliveries (Wang and Toktay, 2008), or to obtain the optimal outbound shipment times (Higginson and Bookbinder, 1995). In constrast, OWMR distribution systems may be used when applying heuristic inventory replenishment or shipment policies, with the objective of optimizing its policy parameters (Marklund, 2011). These simplified models are valuable tools for understanding and addressing complex inventory control challenges.

1.1.2 Supply Lead Times in Inventory Systems

The supply (or replenishment) lead time defines the time from placing an order until the order is available at the ordering stock point. It includes several components as the time for production, transportation, and material handling activities such as picking, loading, receiving, and inspecting. Furthermore, it also includes the waiting time due to delays incurred at the preceding stock point. It may happen that the preceding stock point lacks inventory to fulfill an order immediately. This order is then backlogged until the inventory at the preceding stock point is replenished, leading to waiting times due to stockouts. Moreover, it may also happen that a specific shipment strategy is applied at the preceding stock point, such as consolidating orders or delivering orders only at certain times.

When considering an OWMR distribution system, it becomes feasible to explicitly model waiting times due to stockouts or shipment consolidation strategies at the retailers. Additionally, a deterministic time for transportation and material handling from the warehouse to the retailers may be added. For the first stock point within any inventory system, it is common to aggregate transportation, material handling and waiting times into a constant supply lead time.

1.1.3 Review Periods and Demand Modeling in Inventory Systems

Inventory systems can be divided into periodic review and continuous review systems. If the stock levels are monitored periodically, for instance, once a day or once a week, replenishment orders can only be placed at such a review period. On the other hand, stock levels can also be monitored continuously so that, if necessary, a replenishment order can be placed immediately after the arrival of customer demand.

In general, a periodic review inventory system with a short review period closely approximates the continuous case. Nevertheless, both review approaches offer their respective advantages. Periodic review benefits from coordinating order placements, especially when the demand rate is high. However, in periodic review systems, the safety stock must hedge against demand uncertainties not only during the supply lead time but also during the review period itself. Therefore, safety stocks tend to be higher in periodic review systems compared to continuous review systems. Modern information technologies enable access to real-time data, which has massively reduced monitoring costs in continuous review systems. The advantage of coordinated order placement in a periodic review system can be transferred to a continuous review system by implementing strategies that enable the integration of coordinated dispatches. For instance, one strategy is time-based shipment consolidation, as discussed in Marklund (2011). To summarize, periodic review systems are typically used when the demand rate is high, while continuous review systems are preferred when the demand rate is low.

In practical scenarios, inventory systems usually deal with stochastic customer demand, which is the focus of this work. The way of modeling the customer demand of an inventory system depends on the review period.

In periodic review inventory systems, the demand during the review period is typically modeled as a continuous random variable, even when dealing with a discrete demand pattern. For items with high demand, continuous distributions such as a normal or gamma distribution closely approximate the discrete demand. Even though the normal distribution has been applied in practice and research for a long time, it is worth noting that there remains a small probability of negative demand occurrences. However, if the coefficient of variation is at most 0.3, the probability of negative demand is 0.04 % and can be neglected (Thonemann, 2015). If the coefficient of variation exceeds 0.3, it is advisable to use a distribution where only non-negative demand is possible. An example is the gamma distribution. In case of a low demand items, demand can be modeled using a discrete distribution, such as a Poisson distribution.

In continuous review inventory systems, the timing of customer arrivals is of importance. Therefore, the focus lies on modeling the arrival process of the customer demand. The inter-arrival times between consecutive customers are modeled as stochastic variables, which are assumed to be independent of each other. Such a stochastic process is known as the renewal process. If the inter-arrival times can be described by identically, independently and exponentially distributed random variables and exactly one demand occurs at a time, we define a (pure) Poisson process. A generalization of the Poisson process is the compound Poisson process, where the number of customer demand follows an arbitrary discrete distribution and independence of demand quantities is assumed. The characteristics of the Poisson process offer several analytical advantages, which is why it is often applied in literature when low-demand items are considered.

1.1.4 Replenishment Policies in Inventory Systems

The inflow of goods into a stock point must be managed to ensure that customer demand can be satisfied. Therefore, determining when and how many units should be replenished is of immense importance. However, such a replenishment decision cannot be made based on the physical stock on hand available at the stock point. It is also necessary to consider *outstanding orders* - replenishment orders that have been placed but have not yet arrived at the stock point. Furthermore, a stock point may run out of inventory and, therefore, cannot satisfy customer demand immediately. Assuming that this demand is not lost, it is backordered until a sufficient quantity of units has been replenished. These so-called *backorders* also influence the replenishment decision at the stock point. In summary, stock points base their replenishment decisions on the *inventory position*, which is defined by *stock on hand* plus *outstanding orders* minus *backorders*.

We further define the *inventory level* as *stock on hand* minus *backorders* to facilitate the calculation of holding and backorder costs within an inventory system. Holding costs arise if the inventory level is positive, meaning physical stock is on hand $(IL^+ = \max(IL, 0))$. In contrast, backorder costs are incurred if the inventory is negative, indicating the presence of unfulfilled customer demand $(IL^- = \max(-IL, 0))$.

Two frequently applied replenishment strategies are the (R, Q) and the (s, S) policies. When applying the (R, Q) policy, a replenishment order of size Q is triggered whenever

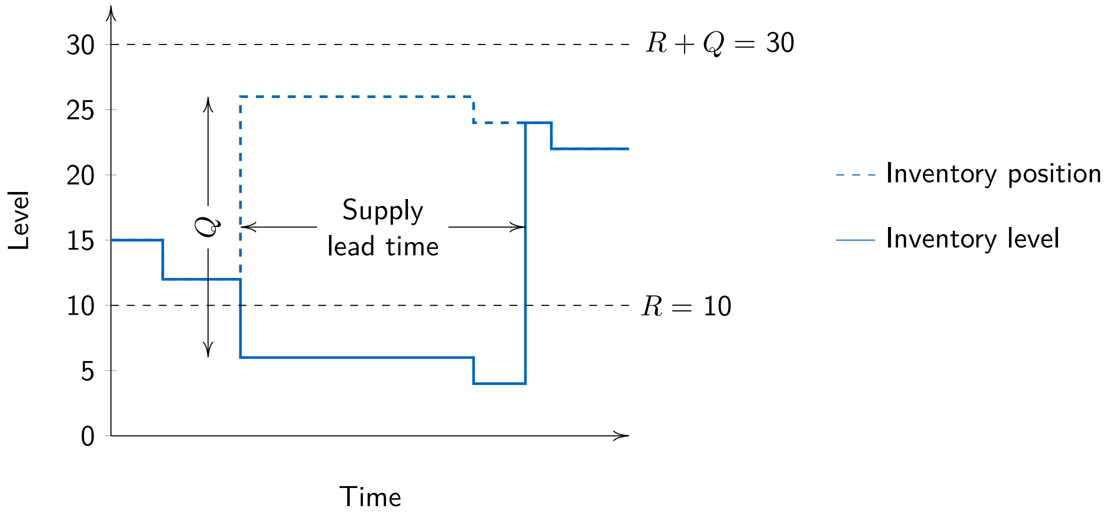

Figure 1.6: Continuous review (R, Q) replenishment policy

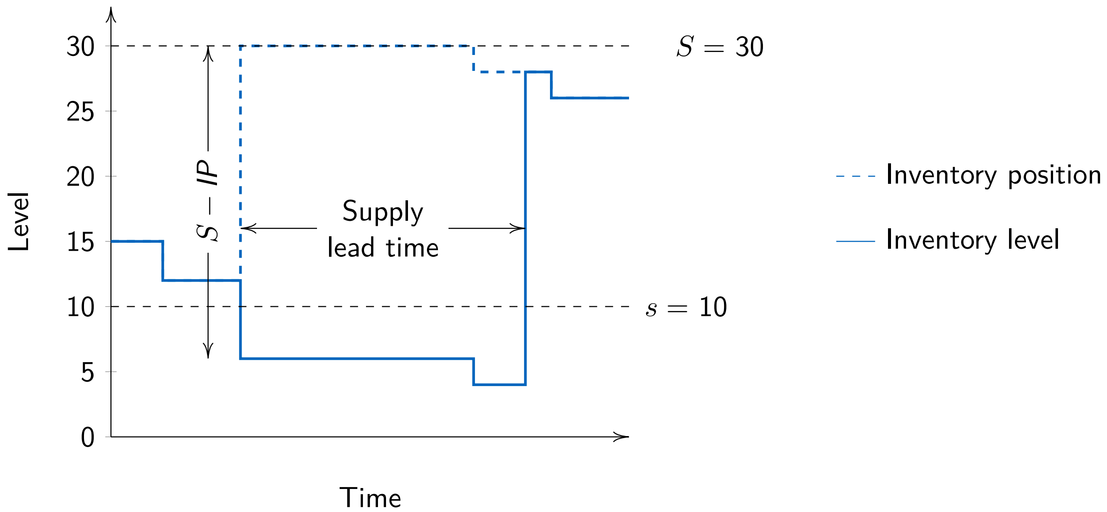

Figure 1.7: Continuous review (s, S) replenishment policy

the inventory position IP falls below or equals the reorder level R $(IP \leq R)$. If the demand size is consistently equal to one, the inventory position exactly hits the reorder level. Consequently, in a continuous review inventory system with order size one, the inventory position will always reach $R + Q$ after placing a replenishment order. An example for a continuous review (R, Q) with random demand, a positive supply lead time, reorder level $R = 10$ and order quantity $Q = 20$ is shown in Figure 1.6.

The main difference between the (R, Q) and the (s, S) policies is the order quantity. When applying the (s, S) policy, an order is placed as soon as the inventory position IP reaches the reorder level s. The inventory position is then raised to S, why the order quantity is changing and equal to $S - IP$. An example for a continuous review (s, S) with random demand, a positive supply lead time, reorder level $s = 10$ and base-stock $S = 30$ is shown in Figure 1.7.

Both policies are equivalent if an inventory system with continuous review and a demand process with demand quantity equal to one is considered. A special case of both policies is the *base-stock policy* or *one-for-one replenishment policy* with base-stock level S. If at least one unit is demanded, an order is placed and the inventory position is raised to S. This policy can be expressed by using either the (s, S) policy with $s = S - 1$ or the (R, Q) policy with $R = S - 1$ and $Q = 1$.

The optimality of the (s, S) policy is proven for a single-echelon inventory system in which holding, backorder and, ordering costs are considered (Scarf et al., 1960; Iglehart, 1963; Veinott, 1966).

1.1.5 Dispatch and Allocation Policies in Inventory Systems

Most conventional inventory models assume that demanded units are shipped to downstream stock points immediately if sufficient inventory is available. In cases of stockouts, demanded units are satisfied as soon as possible after a replenishment order arrives, resulting in *partial deliveries*. Consolidating demands can be effectively integrated through strategically implementing replenishment policies like the (R, Q) or the (s, S) policy at downstream stock points. Nevertheless, partial deliveries may occur due to the model assumptions, and consolidating demands across several stock points is not possible because replenishment decisions are made per stock point.

By contrast, the integration of shipment policies at upstream stock points can handle the consolidation of customer demand across several downstream stock points and additionally the dispatch time, meaning that partial deliveries between regular shipments are not allowed. Higginson and Bookbinder (1994) define the *time-based*, the *quantity-based* and the *hybrid time-and-quantity-based* shipment consolidation policies. The time-based shipping strategy consolidates demand of downstream stock points based on a fixed shipment interval T_c. Shipments to the stock points leave every T_c time units. Random customer demand results in a random shipment quantity at each shipment time. This strategy is widely adopted in industry due to its advantages in scheduling, administration and coordination of processes (Marklund, 2011). When applying the quantity-based policy, the upstream stock point consolidates customer demand until a fixed consolidation quantity of Q_c units is requested and available at the stock point. The consolidation quantity Q_c may be related to a full truckload or container. Therefore, this strategy leads to a high utilization of the transportation mean. However, the time between two consecutive shipments is random why long waiting times for the customer may occur. The hybrid policy is a combination of the time-based and quantity-based policy, meaning that shipments are dispatched when the end of the shipment interval T_c is reached or the quantity of Q_c units has been consolidated.

When sufficient inventory is available at an upstream stock point when a shipment is dispatched, all customer demand is satisfied immediately. However, in cases where the upstream stock point runs out of stock at the moment of shipment, the question arises of how to allocate units to the downstream stock points after a replenishment order has arrived. A widespread and straightforward allocation decision rule commonly found in literature and practice is the first-come-first-serve (FCFS) policy. This policy ensures that backorders are fulfilled in the sequence of order arrival. FCFS is considered an equitable approach and offers analytical advantages, particularly in the field of exact analysis of continuous review inventory systems.

Nonetheless, the FCFS allocation policy may not be optimal, as it does not take into account real-time information regarding the shipment timing, inventory levels at downstream stock points and pipeline inventory. Instead of relying on the straightforward FCFS rule, the upstream stock point may solve an optimization problem each time a shipment is dispatched to the downstream stock points. Since such a decision depends on several aspects, no general optimal allocation policy has been found yet. However, especially in the context of time-based shipment consolidation, Howard and Marklund (2011) and Howard (2013) demonstrate that the FCFS policy performs very well in most cases.

1.1.6 Demand Lead Times in Inventory Systems

In contrast to the supply lead time, which delays the arrival of a replenishment after ordering, the *demand lead time* gives information about demand in advance. The demand lead time is one way of modeling ADI and was first introduced in Hariharan and Zipkin (1995). At the time point where an order occurs, the inventory system gets information about future demand. The due date of this order is exactly after the positive demand lead time. If the *order* is due, we call it *demand* and it is then satisfied directly provided stock on hand is available. In Figure 1.8, we show a representation of a positive and constant demand lead time and the corresponding terminologies.

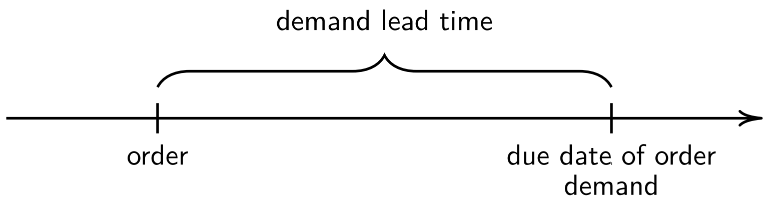

Figure 1.8: Graphical representation of the demand lead time

In a system where the demand lead time is equal to or greater than the supply lead time, the inventory system is able to order units just in time because there is no demand uncertainty. In this case, all demand can be met immediately without the need of keeping inventory. If the demand lead time is smaller than the supply lead time, the demand uncertainty during the supply lead time is reduced by the demand lead time. Therefore, Hariharan and Zipkin (1995) conclude that in a conventional inventory system, a reduction of a constant supply lead time has the same effect as an increase of a constant demand lead time.

When considering a system with a positive demand lead time, the inventory position has to be adapted accordingly. In the moment of customer order arrival, a unit is usually directly allocated to this order even if it is satisfied in the future. Therefore, Gallego and Özer (2001) first define the *modified inventory position* which is equal to *stock on hand* plus *outstanding orders* minus *backorders* minus *observed orders*. Replenishment orders are then placed based on the modified inventory position.

Instead of adhering to the conventional assumption that orders can only be fulfilled once their due date has been reached, Wang and Toktay (2008) propose the concept of *flexible deliveries*. This approach allows orders to be fulfilled immediately after they have occurred.

1.1.7 Performance of Inventory Systems

Two approaches are typically applied to measure the performance of a conventional inventory system that allows unsatisfied demand to be backordered: The *cost-oriented* approach and the *service-oriented* approach. The service-oriented approach aims to minimize the sum of holding costs and ordering costs within the system while adhering to a service constraint. A commonly used service level is the fill rate, which indicates the proportion of demand that can be fulfilled directly from physical stock. In this work, we focus on the cost-oriented approach, with the objective of minimizing the sum of the cost associated with the considered inventory system.

Keeping inventory on stock incurs costs for several reasons. The major part of the holding costs represent opportunity costs for tied up capital. Additionally, material handling, storage, damage and obsolesce, insurance and taxes are included. In summary, holding costs include all costs that vary with the inventory level and are related to warehousing. For instance, if the warehouse space is rented, and more stock would lead to more rented warehouse space, these costs need to be included. The holding costs occur per unit and time unit.

In conventional inventory systems, ordering costs are associated with fixed costs of a replenishment order. These costs may include expenses for transportation, material handling, and administrative tasks. Consequently, considering these costs leads to batch ordering to reduce the total cost of the inventory system. However, in this work, we study the integration of shipment policies in inventory systems to consolidate the demand of downstream stock points rather than relying on batch ordering policies. Therefore, we do not consider ordering costs; instead, fixed and variable shipment costs associated with transportation and material handling are included.

Backorder costs occur if a demand cannot be satisfied from the physical inventory immediately. In such a case, additional costs for administrative tasks, price discounts, and material handling arise. Backorder costs are usually measured per unit and time unit; however, the main challenge lies in determining the cost level.

1.2 Contribution, Methodology and Research Questions

While there is an extensive body of literature addressing inventory control with integrated shipment decisions, the majority focuses on single-echelon inventory systems, such as those by Çetinkaya and Lee (2000) and Çetinkaya et al. (2008). These studies often do not explicitly model inventory levels at customer stock points; instead, they include waiting penalties. Marklund (2011) is the first to explore a time-based consolidation strategy within a multi-echelon distribution system. Furthermore, Malmberg and Marklund (2023) highlight that a quantity-based consolidation strategy results in increased safety stock levels at downstream stock points to hedge against demand uncertainty during extended lead times and show the importance of considering multi-echelon inventory systems.

We aim to contribute to the literature on multi-echelon inventory control by investigating the general case of a hybrid shipment consolidation policy. This shipment strategy offers the advantage of enabling us to analyze under which cost conditions a pure time-based or a pure quantity-based policy could be applied instead, without significant cost increases compared to the dominating hybrid policy. Therefore, our first question is as follows:

> **Research Question 1:** *Under which conditions could a multi-echelon inventory distribution system apply a pure time-based or a pure quantity-based instead of using the dominant hybrid consolidation policy?*

In Chapter 3, we study an inventory system consisting of one warehouse and several retailers, each of the latter receiving stochastic demand from external customers. The warehouse replenishes the retailers' stock according to a time-and-quantity-based shipment consolidation policy, meaning that a shipment is dispatched either when a time-based shipment day is reached or when a consolidation quantity is achieved. Inventory, backorder, and shipment costs are considered in this distribution system. Our main contribution is, on the one hand, the derivation of a new mathematical model which is based on probability theory and aims to compute the exact total cost of the inventory system. After optimizing inventory levels and shipment policy parameters of the distribution system across different instances, we can address the first research question. This chapter is based on Malmberg et al. (2024).

The three consolidation policies presented by Higginson and Bookbinder (1994) are applied both in the literature and in practice. However, despite their widespread application, these heuristic policies may not be reasonable if, e.g., ADI is available. When ADI is modeled as a demand lead time, Wang and Toktay (2008) presented the promising option to satisfy orders before the corresponding due date is reached, known as flexible deliveries. However, they do not apply a consolidation scheme for outbound deliveries.

We contribute to the field of inventory control with integrated shipment decisions by considering a positive demand lead time and answering the following questions:

> **Research Question 2:** *What is the value of incorporating ADI and allowing for flexible deliveries in a single-echelon inventory system when satisfying external orders according to a flexible time-based shipment consolidation policy? What is the effect on the optimal inventory and shipment policy parameters?*

In Chapter 4, we consider a single-echelon inventory system with stochastic orders from an external production facility which are due after a positive demand lead time. The time-based consolidation strategy is adjusted so that orders can be taken ahead of the due date if there is remaining transportation capacity available. We derive a new mathematical model based on probability theory to calculate the total costs of the system. The shipment consolidation policy leads to delayed deliveries to the production facility, who must keep more safety stocks to maintain the same service level. Therefore, late-delivery costs are incurred when a unit is dispatched after its corresponding due date. In contrast, the production facility has to hold stock in advance if units are dispatched before the due date is reached. For this reason, we consider early-delivery costs for all units shipped ahead of its corresponding due date. After optimizing the inventory levels and the shipment interval of the inventory system for numerous instances, we can answer the second research question. This chapter is based on Ralfs and Kiesmüller (2022).

However, after using a heuristic flexible time-based consolidation policy, the question arises whether the proposed shipment policy is a reasonable choice under various cost conditions. Therefore the following question arises:

> **Research Question 3:** *What is the (near-)optimal structure of the outbound shipment policy under time-based dispatching in inventory systems considering ADI and flexible deliveries?*

In Chapter 5, we show how the structural properties of the near-optimal shipment policy for a single-echelon inventory system when external orders are dispatched on a time-based schedule. Therefore, we model the system as a Markov decision process (MDP), where an inventory manager can decide about the consolidation quantity on each shipment day. In order to find the near-optimal shipment policy excluding the influence of backorders at the warehouse, ample stock is assumed at the warehouse. Late and early deliveries will be penalized with costs. Small-scale instances can be solved exactly using value iteration (VI); however large-scale instances cannot be solved with VI because the state space explodes. Therefore, we develop a deep reinforcement learning (DRL) algorithm to estimate the value of the post-decision state. After obtaining the near-optimal outbound shipment policy for numerous large-instances, we can answer the third research question. This chapter is based on Ralfs et al. (2024).

To conclude, Chapter 6 summarizes the main insights of our research. Additionally, we discuss limitations of our models and highlight opportunities for future research.

Chapter

2 Literature Review

The foundation of this work lies in extensive research within the field of stochastic inventory management, with a particular focus on shipment consolidation and ADI. So far, these areas have been examined independently within the existing literature. In the field of inventory management with integrated shipment decisions, demand is due directly. Conversely, in literature focusing on inventory management with positive demand lead times, dispatches take place immediately without consolidating demand. Therefore, the literature relevant to these topics is discussed separately in this work. Key research findings, demonstrating relevance to addressing the research questions, are integrated accordingly. Furthermore, the concluding chapter of the literature review provides a brief overview of literature relating to DRL methods applied to inventory problems.

2.1 Inventory Management with Shipment Consolidation

The three main applied shipment consolidation policies, as detailed in Section 1.1.5, are the time-based, the quantity-based and the hybrid time-and-quantity-based policies. Higginson and Bookbinder (1994) consider a single-echelon inventory system with Poisson demand, linear inventory and fixed shipment costs. They assume ample stock and focus on the outbound dispatch. Through discrete-event simulation, they determine the mean cost per unit and the mean order delay under all three policies, and compare various shipment intervals and consolidation quantities, but they do not optimize these policy parameters.

Further literature on heuristic shipment policies under stochastic demand, although not explicitly integrating inventory decisions, can be found, e.g., in Çetinkaya and Bookbinder

(2003), Bookbinder and Higginson (2002), Mutlu et al. (2010), and Çetinkaya et al. (2014). The former examine the two pure shipment consolidation schemes under private and common carriage, deriving expressions for obtaining the shipment policy parameters under time-based and quantity-based dispatch strategies when inventory and shipment costs are incurred. Inventory holding costs arise from the moment of order arrival until dispatch, while the structure of the shipment costs depends on the type of carriage. Private carriage cause total shipment costs comprising fixed and linear cost components, whereas common carriage only entails weight-dependent linear shipment costs. Bookbinder and Higginson (2002) and Mutlu et al. (2010) focus on hybrid shipment consolidation, offering analytical models to determine the optimal shipment policy parameters T_c and Q_c. The latter study provides analytical expressions for the total cost, enabling a comparison of the three heuristic shipment policies, with $T_c \to \infty$ and $Q_c \to \infty$ denoting the pure quantity-based and the pure time-based policy, respectively. Their findings show that the hybrid policy with $T_c \to \infty$ is preferable in terms of cost. Finally, Çetinkaya et al. (2014) derive expressions to obtain the probability distributions of the maximum and average waiting time of an order for all three shipment policies.

Çetinkaya and Lee (2000) present an analytical model for a stochastic single-echelon inventory system with Poisson demand and time-based shipment consolidation. The stock point applies an (s, S) replenishment policy with reorder level $s = -1$, and a zero supply lead time is assumed, guaranteeing sufficient stock on hand at the moment of dispatch. The total cost consists of linear inventory, linear waiting, linear and fixed replenishment, and linear and fixed dispatching costs. Waiting costs represent the delay in satisfying customer demand, resulting in a penalty for the loss of goodwill. They develop an approximate method to determine near-optimal inventory and shipment policy parameters. In contrast, Axsäter (2001) derives an exact method to obtain the optimal base-stock level S and shipment interval.

For this basic inventory system, several extensions have been published. Çetinkaya et al. (2006) derive close-form expressions for the total cost under a quantity-based consolidation scheme, enabling the optimization of base-stock level S and the shipment quantity. Additionally, they compare the time-based and quantity-based strategies under optimal inventory and shipment decisions, finding that the quantity-based scheme outperforms the time-based scheme in terms of total costs. However, the time-based policy specifies an upper bound on the customers' maximum waiting time. They further suggest hybrid policies with different approximate inventory and shipment policy parameters by combining the parameter values of the two optimal pure policies and compute the total cost by simulation. Çetinkaya et al. (2008) provide easy-to-compute approximations for the base-stock level S and the shipment quantity for the same inventory system assuming a compound renewal process and quantity-based shipment consolidation, highlighting significant cost savings through shipment consolidation. Finally, Wei et al. (2023) present analytical expressions for the hybrid policy, including the special cases of a pure time-based and pure quantity-based policy, for the same single-echelon inventory system with Poisson demand. They compare all three policy under a fixed consolidation cycle length and a fixed replenishment cycle length, without optimization of inventory and shipment decisions. In these circumstances, the pure quantity-based policy outperforms the other policies in average total cost, average order delay and average inventory rate. However, the other policies limit the maximum waiting time.

Previously mentioned papers assume private carriage, while Mutlu and Çetinkaya (2010) investigate the common carriage case with quantity discounts for dispatches for the same basic inventory system as introduced in Çetinkaya and Lee (2000). Quantity discounts result in non-linear shipment costs, motivating the development of search algorithms to determine the optimal inventory and shipment policy parameters for the pure time-based and pure quantity-based strategies.

Chen et al. (2005) adapt the problem setting in Çetinkaya and Lee (2000) by assuming an (R, Q) replenishment policy in a single-echelon inventory system. Moreover, they contribute by not constraining the reorder level to be $R = -1$, but optimizing it, allowing backorders to occur if the inventory position is negative but still above the reorder level R. They derive simple and exact solution approaches to obtain the expected total cost per time unit and identify the optimal inventory and shipment policy parameters under time-based and quantity-based shipment consolidation. Similar to aforementioned literature, they find that the quantity-based strategy can outperform the time-based strategy in term of total cost, while the reverse is not true.

There is also a recently growing body of literature on multi-echelon inventory systems and shipment consolidation. Marklund (2011) considers a stochastic divergent inventory system comprising one warehouse and several groups of retailers, combined with a time-based consolidation scheme, Poisson arrivals, and FCFS allocation. The warehouse manages the inventory using an (R, Q) policy, while retailers apply base-stock policies. He develops an exact recursive evaluation procedure and applies a bounded enumeration to determine the optimal inventory and shipment policy parameters by minimizing the expected total system cost, which includes linear shipment, inventory and backorder costs.

Additionally, two simple yet effective heuristics are presented to compute near-optimal shipment intervals.

Several extensions have been explored in the literature. In contrast to Marklund (2011), who considers the computationally attractive FCFS policy at the warehouse, Howard and Marklund (2011) investigate more general allocation policies. They propose two state-dependent myopic allocation rules and evaluate their cost performance using simulation. With the first policy, allocation is postponed until the moment of shipment, while with the second policy, allocation takes place at the moment of delivery to a group of retailers. Although the state-dependent myopic policies outperform the FCFS policy in special cases, they conclude that the FCFS allocation is a reasonable policy due to its simple implementation.

Stenius et al. (2016) and Johansson et al. (2020) both explore OWMR inventory systems with time-based shipment consolidation and compound Poisson arrivals. The main contribution of Stenius et al. (2016) is the derivation of exact probability distribution for inventory levels at all retailers, enabling the computation of the expected total cost of the system. In contrast, Johansson et al. (2020) consider the same type of inventory system, but develop several approximation methods to mitigate computational complexity.

Stenius et al. (2018) extend the basic OWMR distribution system with time-based shipment consolidation by incorporating non-linear shipment costs. They derive the probability mass function (PMF) for the shipment quantity to a group of retailers, enabling the analysis of any shipment costs structure that is based on the shipment quantity. Sonntag et al. (2023) investigate an inventory routing problem with time-based shipment consolidation and demonstrate how retailer groups should be formed.

Kiesmüller and de Kok (2005) and Malmberg and Marklund (2023) investigate the trade-off between quantity-based shipment consolidation policies and inventory levels in multi-echelon distribution systems. The former provides approximations for computing policy parameters under a given fill-rate, while the latter offers an exact method for evaluating the total system cost subject to a fill-rate constraint. They show that simultaneously optimizing inventory and shipment decisions is highly important, as increasing consolidation quantities result in a significant increase in base-stock levels at the retailers but only a slight decrease in the inventory level at the warehouse.

The literature discussed above examines heuristic consolidation policies and optimizes their policy parameters. Contrary, Higginson and Bookbinder (1995) model a MDP to analyze the optimal outbound shipment quantities for common and private carriage under the assumption of ample stock at the single stock point. They restrict the action space to include only two options: *wait* or *ship*. Inventory holding costs are charged when no shipment is dispatched and orders are further consolidated, while shipment costs arise otherwise. By incorporating a minimum weight required to receive a volume discount on shipments to reflect common carriage, they observe that the optimal policy follows a triangular pattern. This pattern indicates that shipments are dispatched only if the accumulated load is below a first threshold or above a second higher threshold. If the accumulated load is below the first threshold, the time required to accumulate a sufficient quantity may be lengthy, making it favorable to send these units to customers immediately. Conversely, if the accumulated load exceeds the second threshold, holding costs overweight shipment costs, resulting in a dispatch of a shipment. Under private carriage, only constant and fixed shipment costs are considered, and a maximum shipment capacity is assumed, which cannot be exceeded. Therefore, a control-limit policy is optimal, dispatching a shipment once a certain limit is reached. However, due to the curse of a large state space, Higginson and Bookbinder (1995) propose aggregating the states of the MDP to efficiently solve the problem.

Papadaki and Powell (2003) investigate a similar dispatch problem assuming ample stock at the stock point for a multi-item setting. They consider individual linear inventory holding costs for the items awaiting for being dispatched to customers and fixed shipment costs for vehicle dispatch with limited capacity. Items of the same class are shipped according to the FCFS principle. While they also model the problem as a discrete-time MDP, they face challenges due to the multidimensionality of states, outcomes and actions. Therefore, they approximate the value function using an adaptive dynamic programming algorithm to observe the optimal dispatch policy. Similar to Higginson and Bookbinder (1995), they conclude that applying a control-limit policy is optimal, prioritizing customer orders based on their holding costs in descending sequence. A comparable shipment consolidation problem involving two heterogeneous customers with distinct waiting costs are explored in Satır et al. (2018). Additionally, they focus on the uncapacitated problem by modeling it as a continuous-time MDP. Their finding reveals that the optimal policy is a control-limit policy with linear-stepwise thresholds.

The previous literature overview highlights the potential advantages of consolidating shipments, summarizes comparisons of the heuristic policies, and underscores the importance of exploring inventory models that integrate shipment decisions. This thesis analyzes a multi-echelon inventory with hybrid time-and-quantity-based shipment consolidation and determines the corresponding optimal inventory and shipment decisions. Through this exploration, we broaden the scope of findings and yield valuable insights into heuristic shipment policies that are preferable across various scenarios. As there is no research on the joint consideration of shipment consolidation with ADI, this thesis also investigates inventory systems that integrate a time-based consolidation strategy alongside a positive demand lead time, guiding to the discussion about literature on inventory management with ADI.

2.2 Inventory Management with Advance Demand Information

ADI can be modeled in different ways, for example, with dynamic forecast updates, as shown in Toktay and Wein (2001), Schoenmeyr and Graves (2009), and Papier (2016). However, in this work, we model ADI with a positive demand lead time, thus focusing specifically on this aspect within the literature review.

Hariharan and Zipkin (1995) were the first to model ADI using a demand lead time for a single-echelon continuous review inventory system. They assume perfect ADI, meaning that neither the size or the time of demand can be changed, nor can a demand be canceled. Demand arrives according to a Poisson process, is satisfied according to the FCFS policy, and stockouts are backordered. The total cost in the system consists of inventory holding and backorder costs. Based on probability theory, they find that increasing the demand lead time has the same effect as decreasing the supply lead time, effectively shortening the effective supply lead time. As a result, safety stocks can be reduced. If the supply lead time is smaller than the demand lead time, a buy-to-order situation occurs, and no stock keeping is required. Furthermore, they prove a base-stock policy to be optimal for this conventional inventory system without fixed ordering costs.

Following this initial research, a considerable body of literature was developed. Ahmadi et al. (2019b) consider a similar inventory system, but they additionally compensate for ADI by incorporating commitment costs that are strictly increasing with respect to the length of the demand lead time. These costs may represent a bonus for retailers or production facilities accepting a preorder strategy. They optimize the length of the demand lead time and the base-stock level, finding that either a buy-to-stock strategy (standard base-stock policy with no demand lead time) or buy-to-order strategy (no inventory holding since demand lead time equals supply lead time) is optimal. Ahmadi et al. (2019a) find similar results for an assemble-to-order system with two components, while Ahmadi et al. (2020) investigate the same inventory system subject to a time-based service level constraint, revealing that an increase of the demand lead time decreases the average waiting time and increases the service level of a customer. Furthermore, Lu et al. (2003) examine an assemble-to-order system where product orders arrive according to Poisson processes and products may consist of m different components. Components are replenished based on base-stock policies, and the supply lead time is stochastic. They demonstrate that increasing the demand lead time is more effective than decreasing the supply lead time in terms of order fill rates.

In contrast to the the aforementioned literature, Marklund (2006), Du and Larsen (2017), and Figueira et al. (2023) focus on allocation policies when customers order with non-identical demand lead times. Marklund (2006) investigates a continuous-time OWMR distribution system, replenishing stock based on base-stock policies and facing Poisson demand. He proposes a general reservation policy where a reservation time is set for each retailer, along with two simpler policies. One policy assumes that the allocation takes place as soon as the order is known (complete reservation policy), while the other policy assumes that the allocation takes place when the order is dispatched (last-minute reservation policy). He finds that the general policy outperforms the others, and that the last-minute policy prevents incorrect prioritization of customers compared to the compared to the complete reservation policy. Building on the work of Chen (2001), Du and Larsen (2017) include a revenue component that decreases with an increasing demand lead time, and draw similar conclusions as Marklund (2006). Figueira et al. (2023) emphasize that customers should be incentivized to share orders early by receiving a higher individual service. However, late-minute allocation may lead to the opposite effect. Thus, they propose a new strategy based on complete reservation while additionally considering the arrival of the supply pipeline.

In contrast to continuous review inventory systems with ADI, Gallego and Özer (2001) initiated discussions on stochastic single-echelon models with periodic review, non-identical demand lead times, and fixed ordering costs. They were the first to make a replenishment decision based on a *modified inventory position*. The main difference is that a unit is already reserved for one order at the time of arrival, why the classical definition of the inventory position is adapted accordingly (*modified inventory position* equals *stock on hand* plus *outstanding orders* minus *backorders* minus *observed orders*). Moreover, they prove optimality for a state-dependent (s, S) policy.

Several extensions of this initial work have been studied. For instance, Gallego and Özer (2003) investigate a stochastic serial inventory system with periodic review, non-identical demand lead times and a centralized decision maker. They show that a state-dependent base-stock policy is optimal when fixed ordering costs are not included. Özer (2003) examine a divergent inventory system, introduce a heuristic replenishment policy, reaffirming that ADI substitutes inventories or supply lead times. Delving into inventory systems with limited production capacity and ADI, Özer and Wei (2004) find that a state-dependent threshold policy is optimal, indicating that it is optimal either not to produce or produce the full capacity depending on the threshold. Additionally, they conclude that ADI cannot only serve as a substitute for inventory but also for capacity. In a multi-echelon assembly system with ADI, Angelus and Özer (2016) identify the optimal replenishment policy as a state-dependent double-tiered base-stock policy if orders are allowed to expedited through the system for a penalty. Wang and Toktay (2008) were the first to allow orders to be satisfied before the due date, known as flexible deliveries. For identical demand lead times, they prove through dynamic programming that an (s, S) policy is optimal, where the policy parameters depend on the total advance demands. However, in contrast to previous literature demonstrating the equivalence of supply and demand lead times, they illustrate that increasing the demand lead time is preferable to reducing the supply lead time caused by the inclusion of flexible deliveries.

There is also a field of literature considering imperfect ADI, wherein orders can be changed in size or due date, or even cancelled completely. For instance, see Bourland et al. (1996), Thonemann (2002), Tan et al. (2007), Tan et al. (2009), and Topan et al. (2018).

All the literature above demonstrates the significant potential and value of inducing customers to share ADI, as it can reduce uncertainties in the system and, therefore, stock can be reduced without compromising a reduction of customer fill rate. Combining a shipment consolidation policy with ADI and flexible deliveries yields novel insights into logistic processes. This work, therefore, focuses, firstly, on the joint consideration of shipment consolidation and ADI applying a heuristic time-based shipment policy, and secondly, on optimizing the shipment quantities when shipments leave according to time-based intervals and flexible deliveries are allowed.

2.3 Deep Reinforcement Learning in Inventory Control

DRL combines deep learning techniques and reinforcement learning principles, and has developed in recent years into a powerful tool for sequential decision-making. It enables learning to make optimal actions in complex, dynamic, and uncertain environments. Although inventory problems have been analyzed for decades, there are only a small number of problems for which the optimal replenishment policy is known, as analytical derivations are often intractable. For small-scale problems, numerical approaches such as value or policy iteration may solve these problems to optimality. However, these approaches are limited and cannot handle large-scale instances as the dimensions increase. In such cases, DRL is a promising tool to determine structural insights into near-optimal policies for well-known inventory problems. Boute et al. (2022) are convinced that DRL offers significant potential to determine structural insights into near-optimal policies for well-known inventory problems, and provide comprehensive implementation guidelines. To apply DRL methods to inventory problems, it is necessary to model the considered problem as an MDP. As the state and action space expand, for example due to increasing supply lead time, the value or policy function can be efficiently approximated by neural networks to obtain solutions.

Vanvuchelen et al. (2020) analyze the near-optimal policy of the joint replenishment problem using the proximal policy optimization algorithm. This fundamental inventory problem involves jointly replenishing several products from one supplier, while considering a limited truck capacity and inventory, backorders, and ordering costs. They benchmark the DRL policy against heuristic policies and demonstrate that the DRL policy performs closely to optimal in small-scale instances and outperforms the heuristic policies. The proximal policy optimization algorithm is also utilized in van Hezewijk et al. (2023) and performs very well to solve the stochastic capacitated lot sizing problem.

Gijsbrechts et al. (2022) focus on three well-known inventory problems: lost sales, dual sourcing, and multi-echelon inventory problems. For all three models, they employ the asynchronous advantage actor-critic DRL algorithm. In the lost sales model, the near-optimal order quantity needs to be determined considering inventory, lost sales, and ordering costs for a single product. Their algorithm beats most available heuristics, especially with high constant supply lead times. In the dual-source inventory model, replenishments can be made from two sources with different costs and supply lead times, while incurring inventory and backorder costs. The DRL-policy performs very well, achieving an optimality gap of less than 2 %. Lastly, they consider the OWMR inventory models to find the near-optimal replenishment strategy for all locations, accounting for holding, lost sales and expedited delivery costs, and maximum location capacity. Their DRL algorithm surpasses base-stock policies with constant base-stock levels. In conclusion, after initial tuning, the DRL algorithm demonstrates strong performance across all three inventory problems. However, there is a noted lack of structural policy insights.

Oroojlooyjadid et al. (2022) propose a shaped-reward deep Q-network algorithm to obtain near-optimal replenishment quantities in the beer game with a positive and constant supply lead time. The beer game illustrates the bullwhip effect, a phenomenon occurring in a serial supply chain where decentralized orders at each stage lead to amplified fluctuations. De Moor et al. (2022) further enhance the deep Q-network by potential-based reward shaping, and apply this algorithm on a perishable inventory model. Their objective is to identify near-optimal replenishment decisions under stochastic demand and a positive and constant supply lead time. In addition to ordering, holding and shortage costs, costs for perishing are incurred.

All of the previously mentioned methods have in common that they rely on a learned policy network to make decisions. However, in this paper, we present an alternative approach. We focus on estimating the value of the post-decision state, as described in Powell (2022). This method offers the advantage of accommodating a flexible set of feasible decisions across various epochs. By exploiting the value of the post-decision state, we not only facilitate decision-making in the current period but also enable rapid policy evaluation by introducing the current policy. This approach has been shown to be significantly more efficient than learning the values of pre-decision states, which typically requires extensive forward simulations at each epoch.

Chapter 3 Hybrid Shipment Policy in Multi-Echelon Inventory System

The following chapter is based on Malmberg et al. (2024). This chapter analyzes the advantages of shipment consolidation policies in a continuous review multi-echelon inventory system. In Section 3.1, we present a detailed formulation of the investigated inventory system and the hybrid time-and-quantity-based shipment policy, which generalizes the pure time-based and the pure quantity-based policy. Section 3.2 presents the evaluation of the expected total cost for given inventory and shipment parameters. The main contribution of this chapter is the derivation of the PMF of the inventory level at the retailers, which is essential for computing the expected total cost. The evaluation of the inventory level at a specific retailer is presented in Section 3.3. In Section 3.4, we outline the procedure to optimize inventory and shipment decisions simultaneously. Finally, Section 3.5 offers a numerical study that investigates the conditions under which a pure time-based or pure quantity-based policy can be applied instead of the dominating hybrid shipment consolidation policy without significantly increasing the expected total cost. The chapter concludes in Section 3.6 with the main findings and remarks.

3.1 Problem Formulation

We study a centralized continuous review OWMR inventory system where $\mathcal{N}$ non-identical retailers are clustered in $\mathcal{M}$ different retailer groups to which shipments are consolidated. A retailer can only be part of one consolidation group, and Ω^m denotes the set of retailers belonging to retailer group m $(m = 1, 2, \ldots, \mathcal{M})$.

The warehouse uses an (R_0, Q_0) policy to replenish its stock from an outside supplier with a constant supply lead time L_0. This means a replenishment order of Q_0 units is placed when the inventory position (= stock on hand + outstanding replenishment orders − backorders) reaches the reorder level R_0. The order quantity Q_0 is determined by set-up costs and restrictions at the outside supplier. There is free access to point-of-sales data from all the retailers, meaning that there are no economic incentives for the retailers to order in batches. As a consequence, each retailer acts according to a base-stock policy with base-stock level S_i. However, there are fixed costs associated with the material handling and shipping of units between the warehouse and the retailers, reflected in shipment costs.

The retailers face demand that can be modeled as independent Poisson processes, and λ_i denotes the demand rate for retailer i. Consequently, the demand rate of group m is $\lambda^m = \sum_{i \in \Omega^m} \lambda_i$. Moreover, since demand at the retailers is immediately converted to retailer orders, the warehouse faces the superposition of all demand processes, i.e., a Poisson process with rate $\Lambda = \sum_{i=1}^{\mathcal{N}} \lambda_i$.

Complete backordering is assumed at all stock points, and demand is served according to the FCFS policy, meaning that a unit is immediately allocated to satisfy a specific demand.

A distinguishing feature is that the warehouse uses a hybrid time-and-quantity-based shipment consolidation policy to dispatch shipments to retailers in the same consolidation group. This policy combines periodic time-based shipments to consolidation group m every T^m time units with quantity-based shipments of full load carriers of Q^m units in between. We refer to T^m as the *shipment interval* and Q^m as the *consolidation quantity* for retailer group m, and the hybrid shipment policy to retailer group m is defined by the tuple (T^m, Q^m).

A time-based shipment contains all units at the warehouse that are *qualified* for shipment at the time of dispatch. A qualified unit is a unit that has been demanded by a retailer and is available at the warehouse. The shipment intervals may differ for different retailer groups and can, for instance, represent daily or weekly dispatches. The shipment intervals cap the waiting time between shipments to a specific retailer group. In addition, there may be shipments of full quantities, Q^m, dispatched during the shipment intervals if the number of qualified units reaches the consolidation quantity Q^m. This quantity typically represents the size of a load carrier, e.g., a full truckload, a container, or a pallet. The periodic dispatch schedule is not altered or restarted by the occurrence of a quantity-based shipment. Having periodic shipment schedules fixed over a specified planning horizon or constant period is often desirable in practice. However, to optimize system performance, joint consideration of T^m and Q^m is key.

Note that using the hybrid policy (T^m, Q^m) with $Q^m = 1$ corresponds to a system without any shipment consolidation policy. Moreover, if $Q^m \to \infty$, the shipment consolidation policy becomes a pure time-based consolidation policy, and for $T^m \to \infty$, the system becomes a pure quantity-based shipment consolidation policy. Hence, our present work generalizes the existing literature on shipment consolidation in OWMR systems by combining the studied time-based and quantity-based policies into a hybrid policy that dominates the existing policies when optimized.

The number of units in a time-based shipment to a retailer group m, I^m, is a random variable that can never be larger than $Q^m - 1$. Otherwise, a quantity-based shipment of Q^m units would already have been dispatched. Furthermore, the number of quantity-based shipments within a shipment interval T^m is random and denoted by J^m.

The transportation time from the warehouse to retailer i, L_i, is constant and the same for all shipments (time-based and quantity-based shipments). This implies that units shipped to retailer i cannot cross in time. However, transportation times may vary across retailers.

The shipment consolidation policy and the FCFS allocation imply that the stock on hand at the warehouse can be divided into; *(i)* unreserved stock on hand $(IL_0^+ > 0)$ that has not yet been demanded by a retailer, and *(ii)* reserved stock on hand that has qualified for shipment and awaits transport on the next available time-based or quantity-based shipment. The reserved stock on hand at the warehouse for retailer i, denoted by W_i.

If there is no stock available at the warehouse at the moment a retailer places an order, the demand is backordered until the replenishment with the reserved unit arrives. The backorder is then cleared, and the unit is assigned to the reserved stock on hand. The inventory level of the unreserved stock at the warehouse is $IL_0 = IL_0^+ - IL_0^-$, where IL_0^- denotes the backorders. Similarly, the inventory level at retailer i is denoted by $IL_i = IL_i^+ - IL_i^-$, where IL_i^+ denotes the stock on hand and IL_i^- denotes the backorders.

In accordance with the existing literature, we focus on the expected total cost per time unit in the system, TC_M, defined in Equation (3.1). The total cost includes the expected holding costs at the central warehouse and all the retailers, expected backorder costs at all retailers, and expected shipment costs from the warehouse to each retailer in every retailer group. The holding costs rate per unit and time unit at stock point i is denoted by h_i, and p_i denotes the backorder costs rate per unit and time unit at retailer i. The shipment costs consist of fixed costs per time-based shipment to consolidation group m, ω_t^m, and fixed costs for each quantity-based shipment, ω_q^m. Furthermore, there are variable shipment costs of c_t^m per shipped unit in a time-based shipment and c_q^m for each shipped unit in a quantity-based shipment. Note that the variable shipment costs are only relevant for decision-making if $c_t^m \neq c_q^m$.

Figure 3.1 shows an exemplary graphical representation of the distribution system.

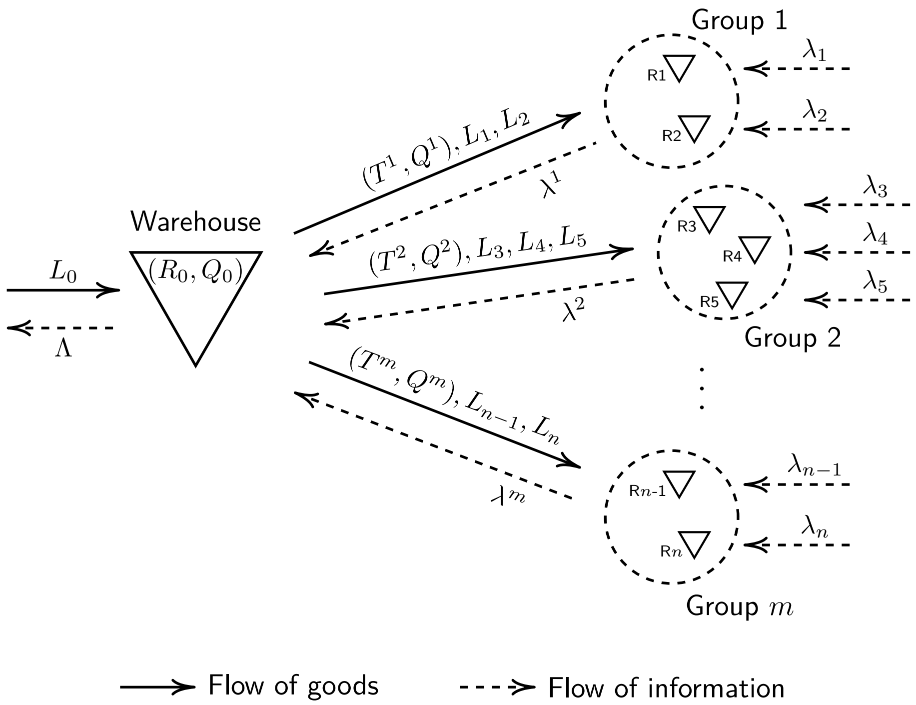

Figure 3.1: Flow of information and goods in the OWMR inventory system

TC_M is a function of the warehouse reorder R_0, the retailers' base-stock levels $\mathbf{S} = (S_1, S_2, \ldots, S_{\mathcal{N}})$, the shipment intervals $\mathbf{T} = (T^1, T^2, \ldots, T^{\mathcal{M}})$, and the consolidation quantities $\mathbf{Q} = (Q^1, Q^2, \ldots, Q^{\mathcal{M}})$ as shown in Equation (3.1).

$$
\begin{aligned}
TC_M(R_0, \mathbf{S}, \mathbf{T}, \mathbf{Q}) = h_0 &\left(E[IL_0^+] + \sum_{i=1}^{\mathcal{N}} E[W_i] \right) \\
&+ \sum_{i=1}^{\mathcal{N}} \left(h_i E[IL_i^+] + p_i E[IL_i^-] \right) \\
&+ \sum_{m=1}^{\mathcal{M}} \frac{1}{T^m} \left(\omega_t^m + \omega_q^m E\left[J^m\right] \right) \\
&+ \sum_{m=1}^{\mathcal{M}} \frac{1}{T^m} \left(c_t^m E\left[I^m\right] + c_q^m E\left[J^m\right] Q^m \right)
\end{aligned} \tag{3.1}
$$

In order to determine the cost-minimizing combination of the decision variables, we formulate the following non-linear optimization problem:

$$
\min\ TC_M(R_0, \mathbf{S}, \mathbf{T}, \mathbf{Q}) \qquad R_0, \mathbf{S} \in \mathbb{Z}, \mathbf{T} \in \mathbb{Q}_{>0}, \mathbf{Q} \in \mathbb{N}_0 \tag{3.2}
$$

3.2 Evaluating of the Total Cost

In this section, we explain how to compute the expected total cost as described in Equation (3.1), assuming the PMFs of all inventory levels, IL_i for $i \in \{1, 2, \ldots, \mathcal{N}\}$, are known. Section 3.3 subsequently presents the method for obtaining these PMFs.

The first two terms of the expected total costs per time unit in Equation (3.1) are the expected holding costs at the warehouse, divided into costs for expected unreserved stock on hand ($h_0 E[IL_0^+]$) and expected reserved stock on hand ($h_0 \sum_{i=1}^{\mathcal{N}} E[W_i]$). It is noteworthy that the unreserved stock on hand at the warehouse and the backorders at the warehouse are not affected by the shipment consolidation policy. The expected unreserved stock on hand can, therefore, be computed in the same way as in a system without shipment consolidation, see Axsäter (2015).

$$
E[IL_0^+] = E[IL_0^-] + R_0 + \frac{Q_0 + 1}{2} - \Lambda L_0 \tag{3.3}
$$

To compute the expected number of backordered units at the warehouse, $E[IL_0^-]$, we use that in a system without shipment consolidation, all units outstanding at retailer i are either backordered at the warehouse or under transport from the warehouse to this retailer. Letting $E[IL_i|Q^m = 1]$ denote the expected inventory level at retailer i in a system without shipment consolidation, the expected number of outstanding orders at retailer i in the system is $S_i - E[IL_i|Q^m = 1]$. Subtracting the expected number of units under transport to retailer i, i.e., $\lambda_i L_i$, renders the expected number of backorders for retailer i at the warehouse. We therefore get

$$E[IL_0^-] = \sum_{i=1}^{\mathcal{N}} (S_i - E[IL_i|Q^m = 1] - \lambda_i L_i) \,. \tag{3.4}$$

The expected reserved stock on hand at the warehouse destined for retailer i, $E[W_i]$, can be obtained from Proposition 1.

Proposition 1. *Let $E[IL_i|Q^m = 1]$ be the expected inventory level for retailer i in group m in a system without shipment consolidation. Furthermore, let $E[IL_i|T^m, Q^m]$ be the expected inventory level at retailer i when using a hybrid shipment consolidation policy with shipment interval T^m and consolidation quantity Q^m.*

The expected reserved stock on hand at the warehouse for retailer i is given by

$$E[W_i] = E[IL_i|Q^m = 1] - E[IL_i|T^m, Q^m]. \tag{3.5}$$

The proof of Proposition 1 is given in Appendix 3.7.1 and follows the same logic as in Malmberg and Marklund (2023) for quantity-based shipment consolidation.

The next set of terms of the expected total cost in Equation (3.1) consists of the summation of the retailers' expected holding costs ($h_i E[IL_i^+]$) and backorder costs ($p_i E[IL_i^-]$). These are directly obtained from the PMF of the inventory level at the respective retailer, thus,

$$E\left[IL_i^+\right] = \sum_{j=0}^{S_i} jPr(IL_i = j) \tag{3.6}$$

$$E\left[IL_i^-\right] = \sum_{j=0}^{\infty} jPr(IL_i = -j). \tag{3.7}$$

The final two terms in Equation (3.1) are the shipment costs for moving items from the warehouse to the retailers. Both the expected number of units in a time-based shipment, $E[I^m]$, and the expected number of quantity-based shipments during the shipment interval T^M, $E[J^m]$, can be derived from the PMF of the number of units that qualify for shipment between two arbitrary time-based shipments to group m, denoted by K^m. It is noteworthy that K^m under the hybrid policy is exactly the same as in a system with a pure time-based shipment consolidation policy. Thus, K^m can be computed using the method presented in Stenius et al. (2018). A brief summary of their approach is provided in Appendix 3.7.2.

Note that every multiple of Q^m units that have been qualified for shipment between two time-based shipments will be part of a quantity-based shipment. Thus, the expected number of quantity-based shipments that occur for each time-based shipment is given by the quotient of the Euclidean division of K^m and Q^m. Thus,

$$E\left[J^m\right] = \sum_{j=0}^{\infty} \left\lfloor \frac{j}{Q^m} \right\rfloor Pr(K^m = j). \tag{3.8}$$

Moreover, the expected number of units shipped on a time-based shipment is given by the remainder of the same division in Equation (3.9). The remainder is defined as $\text{mod}_n\left(x\right) = x - n\lfloor \frac{x}{n} \rfloor$, such that we obtain

$$E\left[I^m\right] = \sum_{j=0}^{\infty} \text{mod}_{Q^m}\left(j\right) Pr(K^m = j). \tag{3.9}$$

3.3 Evaluating the Inventory Level at the Retailers

In this section, we present the method for obtaining the PMF of the inventory level at each retailer. We focus on retailer $\mathcal{N}$ in consolidation group $\mathcal{M}$, but the same methodology applies to any retailer within any consolidation group.

The approach is based on considering the system at an arbitrary time, τ, just before a customer arrives at retailer $\mathcal{N}$. Based on the PASTA (Poisson arrivals see time averages) property of Poisson arrivals (Wolff, 1982), the arriving customers see the steady-state distribution of the inventory level at retailer $\mathcal{N}$, $IL_{\mathcal{N}}$. Because the inventory position, $IP_{\mathcal{N}}$, equals $S_{\mathcal{N}}$ at all times, $IL_{\mathcal{N}}$ is uniquely determined by the outstanding units that have been ordered by retailer $\mathcal{N}$ but have not yet arrived, $O_{\mathcal{N}}$, as $IL_{\mathcal{N}} = S_{\mathcal{N}} - O_{\mathcal{N}}$. Thus, by deriving the PMF of the number of outstanding units for retailer $\mathcal{N}$ at time τ, we get the PMF of $IL_{\mathcal{N}}$ for any base-stock level $S_{\mathcal{N}}$.

To simplify the analysis, we divide the total number of outstanding units for retailer $\mathcal{N}$ at time τ into units ordered before time $\tau - L_{\mathcal{N}}$, denoted by $\widehat{O}_{\mathcal{N}}$, and units ordered after time $\tau - L_{\mathcal{N}}$, denoted by $\widetilde{O}_{\mathcal{N}}$. For a unit to arrive at retailer $\mathcal{N}$ before τ, it must have been dispatched from the warehouse no later than $\tau - L_{\mathcal{N}}$. It follows that any unit ordered after $\tau - L_{\mathcal{N}}$ is outstanding at τ. However, units ordered before $\tau - L_{\mathcal{N}}$ are only outstanding if their shipment has been delayed until after $\tau - L_{\mathcal{N}}$. This delay at the warehouse can be caused by a stock-out and/or a shipment delay due to the consolidation policy. $\widetilde{O}_{\mathcal{N}}$ and $\widehat{O}_{\mathcal{N}}$ are independent random variables as they are determined by customer demand in disjoint time intervals, and $O_{\mathcal{N}} = \widetilde{O}_{\mathcal{N}} + \widehat{O}_{\mathcal{N}}$. Consequently, the PMF of $IL_{\mathcal{N}}$ can be obtained from Equation (3.10).

$$\begin{aligned} Pr(IL_{\mathcal{N}} = j) &= Pr(S_{\mathcal{N}} - O_{\mathcal{N}} = j) \\ &= \sum_{x=0}^{S_{\mathcal{N}}-j} Pr(\widetilde{O}_{\mathcal{N}} = S_{\mathcal{N}} - j - x) Pr(\widehat{O}_{\mathcal{N}} = x) \end{aligned} \tag{3.10}$$

By definition, the outstanding units at time τ ordered after $\tau - L_{\mathcal{N}}$, $\widetilde{O}_{\mathcal{N}}$, correspond to the demand during the transportation time $L_{\mathcal{N}}$. As the demand at retailer $\mathcal{N}$ follows a Poisson process with rate $\lambda_{\mathcal{N}}$, we know that $\widetilde{O}_{\mathcal{N}}$ is Poisson distributed with parameter $\lambda_{\mathcal{N}} L_{\mathcal{N}}$.

The distribution of $\widehat{O}_{\mathcal{N}}$ is more challenging to obtain as it directly depends on the stock availability at the warehouse and the shipment consolidation policy. To derive this distribution, we first let $\widehat{O}^{\mathcal{M}}$ denote the total number of outstanding units for the retailers in group $\mathcal{M}$ (where retailer $\mathcal{N}$ resides) at time τ that have been ordered before $\tau - L_{\mathcal{N}}$. As a consequence of the Poisson demand and FCFS allocation, we can assert that the distribution of $\widehat{O}_{\mathcal{N}}$ can be obtained from a binomial disaggregation of $\widehat{O}^{\mathcal{M}}$, as given in Equation (3.11).

$$Pr(\widehat{O}_{\mathcal{N}} = x | \widehat{O}^{\mathcal{M}} = y) = \binom{y}{x} \left(\frac{\lambda_{\mathcal{N}}}{\lambda^{\mathcal{M}}} \right)^{x} \left(1 - \frac{\lambda_{\mathcal{N}}}{\lambda^{\mathcal{M}}} \right)^{y-x}, \qquad x \leq y \tag{3.11}$$

The PMF of $\widehat{O}_{\mathcal{N}}$ can then easily be obtained from the PMF of $\widehat{O}^{\mathcal{M}}$ by the law of total probability. To facilitate this derivation of $\widehat{O}^{\mathcal{M}}$, we let $\tau' = \tau - L_{\mathcal{N}}$ and focus on what happens in the system before τ' since nothing that happens after τ' affects $\widehat{O}^{\mathcal{M}}$. All units ordered by retailer group $\mathcal{M}$ before τ' can be placed in a priority list with the priorities determined by the FCFS policy. The priority is defined such that the most recently ordered unit before τ' has priority 1 (i.e., lowest priority), the unit ordered before that has priority 2, and so on. The priority list (1,2,3,...) contains all units ordered by retailer group $\mathcal{M}$ before τ'.

The approach we follow is to consider an arbitrary unit with priority n, referred to as the *considered* n^{th} *unit*, and determine the probability that this unit is dispatched to retailer group $\mathcal{M}$ before (or after) τ'. Note that there are $n-1$ units ordered after the considered n^{th} unit and before τ'. Hence, if the considered n^{th} unit is dispatched before τ', there can be at most $n-1$ outstanding units at the warehouse for retailer group $\mathcal{M}$ at time τ'. This implies

$$Pr(\widehat{O}^{\mathcal{M}} \leq n-1) = \\ Pr(\text{the considered } n^{th} \text{ unit has been dispatched before } \tau'). \tag{3.12}$$

It follows that the PMF of $\widehat{O}^{\mathcal{M}}$ can be determined as

$$Pr(\widehat{O}^{\mathcal{M}} = x) = Pr(\widehat{O}^{\mathcal{M}} \leq x) - Pr(\widehat{O}^{\mathcal{M}} \leq x-1). \tag{3.13}$$

Following from Equation (3.10) and Equation (3.11), we can use Equation (3.12) and Equation (3.13) to obtain the PMF of $IL_{\mathcal{N}}$ for any $S_{\mathcal{N}}$. In the following, we focus on determining the probability that the considered n^{th} unit in the priority list has been dispatched from the warehouse to retailer group $\mathcal{M}$ before τ'. In Section 3.3.1, we present the steps in the general method for determining the probabilities $Pr(\widehat{O}^{\mathcal{M}} \leq n-1)$ for all n $(n = 1, 2, 3, \ldots)$ in Equation (3.12). For these steps, we define probabilities, which are explained and derived in Section 3.3.2 and Section 3.3.3.

3.3.1 Determining $Pr(\widehat{O}^{\mathcal{M}} \leq n-1)$

To proceed with the analysis, recall that a qualified unit refers to a unit that is both available at the warehouse and demanded by one of the retailers. In a system without shipment consolidation, qualified units are immediately dispatched from the warehouse. Conversely, in a system with a shipment consolidation policy, qualified units may need to wait at the warehouse before a shipment is dispatched. For the considered hybrid policy, there are two important observations regarding how and when units will be shipped:

1. All units qualified for shipment before the dispatch of a time-based shipment will either be shipped before or with this time-based shipment.

2. If the number of qualified units reaches the consolidation quantity, all of these units are dispatched on a quantity-based shipment the moment the *last unit* in this shipment qualifies for shipment.

Based on these two observations, we can identify two possible events where the considered n^{th} unit is dispatched before τ'. The first possibility is that the considered n^{th} unit is shipped with (or before) the most recent time-based shipment before τ'. This event is referred to as *Shipment Option 1* (*ShipOp1*).

The second possibility is that the considered n^{th} unit is dispatched with a quantity-based shipment after the most recent time-based shipment but before τ'. This event, referred to as *Shipment Option 2* (*ShipOp2*), requires that the considered n^{th} unit qualifies for shipment after the most recent time-based shipment before τ' and that the last unit in the same quantity-based shipment as the considered n^{th} unit qualifies for shipment before τ'. The definitions of these two events for the considered n^{th} unit are summarized below, together with some important time instances that we use in the analysis. See also the illustration in Figure 3.2 and Figure 3.3.

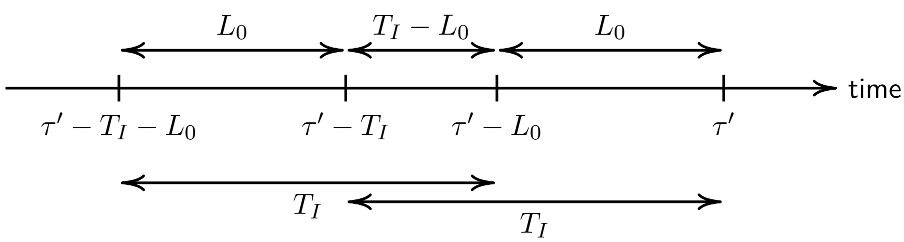

Figure 3.2: Timeline with the four mentioned time instances for Case A, $T_I > L_0$

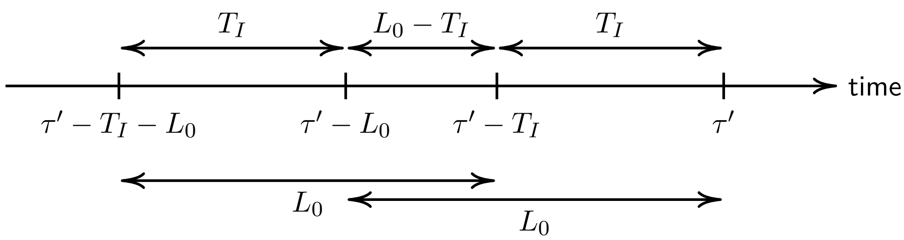

Figure 3.3: Timeline with the four mentioned time instances for Case B, $T_I \leq L_0$

$\tau' - T_I$	The time of dispatch for the most recent time-based shipment to retailer group $\mathcal{M}$ before τ'
$\tau' - L_0$	The latest time an order can be placed by the warehouse so that the units in this order are available at the warehouse at τ'
$\tau' - T_I - L_0$	The latest time an order can be placed by the warehouse so that the units in this order are available at the warehouse at $\tau' - T_I$
ShipOp1	The considered n^{th} unit is either part of the most recent time-based shipment before τ' dispatched at time $\tau' - T_I$, or it belongs to another shipment (time-based or quantity-based shipment) dispatched earlier. This implies that the considered n^{th} unit is qualified for shipment at or before the most recent time-based shipment dispatched before τ'.
ShipOp2	The considered n^{th} unit is part of a quantity-based shipment dispatched after the most recent time-based shipment but before τ'. This can only occur if the last unit in this quantity-based shipment qualifies for shipment prior to τ' and the considered n^{th} unit qualifies for shipment after $\tau' - T_I$.

We can express the probability of interest from Equation (3.12) using the events ShipOp1 and ShipOp2,

$$Pr(\text{the considered } n^{th} \text{ unit is dispatched before } \tau') = Pr\left(\text{ShipOp1} \cup \text{ShipOp2}\right). \tag{3.14}$$

For the considered n^{th} unit to be shipped with a time-based shipment according to ShipOp1, it must qualify for shipment before $\tau' - T_I$. This means it must be demanded by a retailer before $\tau' - T_I$ and ordered by the warehouse before $\tau' - T_I - L_0$. Similarly, for the considered n^{th} unit to be shipped with a quantity-based shipment according to ShipOp2, the last unit in the quantity-based shipment, including the considered n^{th} unit, must qualify for shipment before τ'. Thus, the last unit must be demanded by a retailer before τ' and ordered by the warehouse before $\tau' - L_0$.

By definition, a time-based shipment occurs every $T^{\mathcal{M}}$ time units independently of the stochastic demand. Because the latter follows independent Poisson processes, τ' may occur with equal probability at any time between two consecutive time-based shipments. Consequently, T_I is a uniform random variable on the interval $[0, T^{\mathcal{M}}]$. It follows that in case $T^{\mathcal{M}} > L_0$, it is possible for T_I to be smaller than, equal to, or larger than L_0. This means there are two cases to consider, see Figure 3.2 and Figure 3.3: Case A, where $T_I > L_0$ and, thus, $\tau' - T_I$ occurs before $\tau' - L_0$, and Case B, where $T_I \leq L_0$ and $\tau' - T_I$ occurs after (or at the same time as) $\tau' - L_0$. Moreover, as illustrated in Figure 3.2 and Figure 3.3, the lengths of the intervals also depend on T_I.

If $T^{\mathcal{M}} \leq L_0$, only Case B can occur. If $T^{\mathcal{M}} > L_0$, we know that $T_I \sim Uniform(0, T^{\mathcal{M}})$, where $Uniform$ denotes the uniform distribution, and L_0 is constant, why we have $Pr(\text{Case A}) = Pr(T_I > L_0) = 1 - \frac{L_0}{T^{\mathcal{M}}}$ and $Pr(\text{Case B}) = 1 - Pr(T_I > L_0)$. Moreover, given Case A, we have $T_I \sim Uniform(L_0, T^{\mathcal{M}})$, and given Case B, we have $T_I \sim Uniform\left(0, \min(L_0, T^{\mathcal{M}})\right)$. As the sequence of the events differs between these two cases, they will be treated separately. We may use the law of total probability to rewrite Equation (3.14) as

$$\begin{aligned} &Pr(\text{the considered } n^{th} \text{ unit is dispatched before } \tau') = \\ &\qquad Pr(\text{ShipOp1} \cup \text{ShipOp2} | \text{Case A}) Pr(\text{Case A}) \\ &\qquad + Pr(\text{ShipOp1} \cup \text{ShipOp2} | \text{Case B}) Pr(\text{Case B}). \end{aligned} \tag{3.15}$$

Henceforth, we focus on the conditional probability of ShipOp1 and ShipOp2 given Case A and Case B, respectively.

For Case A and Case B, the defined time instances form three non-overlapping time intervals, see Figure 3.2 and Figure 3.3. These intervals differ for Case A and Case B with respect to the sequence in which the time instances occur. Moreover, the length of the intervals differs with respect to the random variable T_I. However, for the analysis, it is the demand in each of the three time intervals that is of importance rather than the length of the interval. We may consider the demand in each interval to be stochastic variables determined by the (random) length of the interval and the stochastic customer arrival process. When analyzing the outstanding units for retailer group $\mathcal{M}$, we do not have to distinguish between the remaining retailer groups. Only the aggregated demand from these retailer groups influences the number of outstanding units for retailer group $\mathcal{M}$. Therefore, without loss of generality, when considering retailer group $\mathcal{M}$,

Table 3.1: Demand realizations during the three time intervals given Case A and Case B

		Case A	Case B
Retailer group $\mathcal{M}$	Interval 1	$D^{\mathcal{M}}(\tau' - L_0, \tau') = \alpha_1^A$	$D^{\mathcal{M}}(\tau' - T_I, \tau') = \alpha_1^B$
	Interval 2	$D^{\mathcal{M}}(\tau' - T_I, \tau' - L_0) = \alpha_2^A$	$D^{\mathcal{M}}(\tau' - L_0, \tau' - T_I) = \alpha_2^B$
	Interval 3	$D^{\mathcal{M}}(\tau' - T_I - L_0, \tau' - T_I) = \alpha_3^A$	$D^{\mathcal{M}}(\tau' - T_I - L_0, \tau' - L_0) = \alpha_3^B$
Other retailer groups $\overline{\mathcal{M}}$	Interval 1	$D^{\overline{\mathcal{M}}}(\tau' - L_0, \tau') = \beta_1^A$	$D^{\overline{\mathcal{M}}}(\tau' - T_I, \tau') = \beta_1^B$
	Interval 2	$D^{\overline{\mathcal{M}}}(\tau' - T_I, \tau' - L_0) = \beta_2^A$	$D^{\overline{\mathcal{M}}}(\tau' - L_0, \tau' - T_I) = \beta_2^B$
	Interval 3	$D^{\overline{\mathcal{M}}}(\tau' - T_I - L_0, \tau' - T_I) = \beta_3^A$	$D^{\overline{\mathcal{M}}}(\tau' - T_I - L_0, \tau' - L_0) = \beta_3^B$

we will refer to all the remaining retailer groups as group $\overline{\mathcal{M}}$. We let $D^{\mathcal{M}}(s,t)$ denote the random number of demands in time interval $(s,t]$ at retailer group $\mathcal{M}$. Similarly, $D^{\overline{\mathcal{M}}}(s,t)$ denotes the demand that occurs at the other retailer groups in the same time interval.

By conditioning on the demand in each of the three time intervals for both Case A and Case B, we can derive the conditional probability of ShipOp1 and ShipOp2. The demand realizations for each interval are defined in Table 3.1.

Following the notations in Table 3.1, we refer to the most recent time interval before τ' as Interval 1. The second most recent interval is referred to as Interval 2, and the last interval is Interval 3. Figure 3.4 and Figure 3.5 illustrate the three time intervals for Case A and Case B, respectively. These figures also depict two possible demand realizations for retailer group $\mathcal{M}$ and the other retailer groups $\overline{\mathcal{M}}$.

The demand realizations during these intervals are summarized in the vectors $\boldsymbol{\alpha^i} = (\alpha_1^i, \alpha_2^i, \alpha_3^i)$ and $\boldsymbol{\beta^i} = (\beta_1^i, \beta_2^i, \beta_3^i)$ where $i = A$ for Case A and $i = B$ for Case B.

Recall that all units ordered by retailer group $\mathcal{M}$ prior to τ' are enumerated in a priority list. By relating the priority of a unit to the demand realizations from retailer group $\mathcal{M}$, $\boldsymbol{\alpha^i}$, we can determine in which interval the demand for this unit takes place. To exemplify, let us consider Case B (Figure 3.5). If $\alpha_1^B \geq n > 0$, we know that the considered n^{th} unit is ordered in Interval 1 and thus after $\tau' - T_I$, that is, after the dispatch of the most recent time-based shipment before τ'. Therefore, it cannot be part of this shipment.

In addition to the number of demands in the respective time intervals, we also condition the analysis on the number of backorders at the warehouse destined for retailer group $\mathcal{M}$ at time $\tau' - T_I$, denoted by $B_0^{\mathcal{M}}(\tau' - T_I)$. If the considered n^{th} unit is one of these

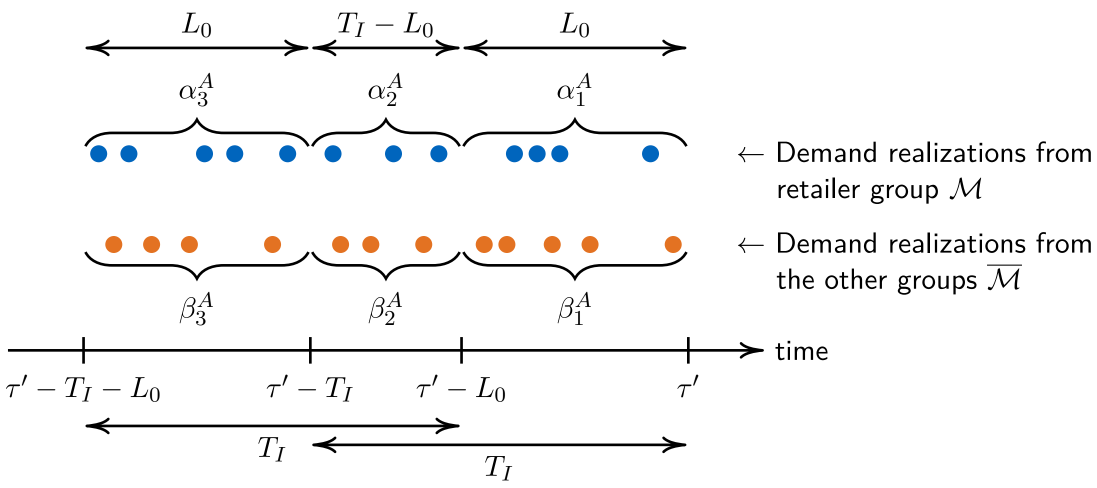

Figure 3.4: Illustration of time instances and events in the corresponding intervals for Case A, i.e., $T_I > L_0$.

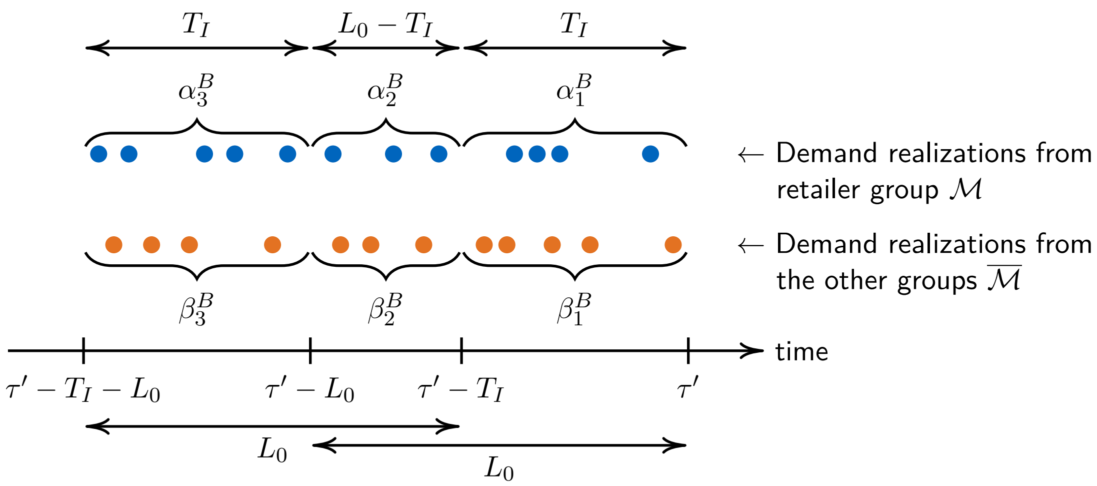

Figure 3.5: Illustration of time instances and events in the corresponding intervals for Case B, i.e., $T_I \leq L_0$.

backordered units, we know that it cannot be part of the time-based shipment at $\tau' - T_I$ as it is not available at the warehouse at this time. Moreover, we condition the analysis on the inventory position at the warehouse at time $\tau' - T_I - L_0$, denoted by $IP_0(\tau' - T_I - L_0)$. The reason is that when considering the last unit in a possible quantity-based shipment containing the considered n^{th} unit, we need to determine the time this last unit was ordered by the warehouse.

We define the realizations' of these random variables as

$$IP_0\left(\tau' - T_I - L_0\right) = ip$$
$$B_0^{\mathcal{M}}\left(\tau' - T_I\right) = b_0^{\mathcal{M}}.$$

Note that the total number of backorders at the warehouse at time $\tau' - T_I$ depends on the inventory position at the warehouse and the demand realizations between $\tau' - T_I - L_0$ and $\tau' - T_I$. These dependencies are dealt with in the derivation of the joint PMF of $\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip$ and $b_0^{\mathcal{M}}$ in Section 3.3.3. For convenience, we will use the notation $\mathcal{P}_i(x) = Pr(X = x|\text{Case } i)$ when it is unambiguously defined from other notations that x is a realization of the random variable X conditioned on Case $i \in \{\text{A,B}\}$. In the same spirit, we will use $\mathcal{P}_i(x|y)$ to denote $Pr(X = x|Y = y, \text{Case } i)$, and $\mathcal{P}_i(x, y)$ to denote $Pr(X = x \text{ and } Y = y|\text{Case } i)$.

Let us define the joint PMF as $\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}})$ for Case $i \in \{\text{A,B}\}$. The conditional probabilities for ShipOp1 ∪ ShipOp2 (i.e., that the considered n^{th} unit is dispatched before or at τ') for Case A and Case B are given by $\mathcal{P}_i\left(\text{ShipOp1}|\boldsymbol{\alpha^i}, b_0^{\mathcal{M}}\right) + \mathcal{P}_i\left(\text{ShipOp2}|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}}\right)$, $i \in \{\text{A,B}\}$, respectively. Note that we condition the analysis on $\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip$, and $b_0^{\mathcal{M}}$. However, for ShipOp1 we get sufficient information from $\boldsymbol{\alpha^i}$ and $b_0^{\mathcal{M}}$; thus, $\mathcal{P}_i\left(\text{ShipOp1}|\boldsymbol{\alpha^i}, b_0^{\mathcal{M}}\right) = \mathcal{P}_i\left(\text{ShipOp1}|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}}\right)$.

Using these probabilities, we can apply the law of total probability to derive an expression for the unconditional probability of the considered n^{th} unit being dispatched before τ'.

$$\begin{aligned}
&Pr(\text{the considered } n^{th} \text{ unit is dispatched before } \tau') = \\
&\qquad Pr(\text{Case A}) \sum_{(\boldsymbol{\alpha^A}, \boldsymbol{\beta^A}, ip, b_0^{\mathcal{M}}) \in \Theta_A} \Big(\mathcal{P}_A\left(\text{ShipOp1}|\boldsymbol{\alpha^A}, b_0^{\mathcal{M}}\right) \\
&\qquad + \mathcal{P}_A\left(\text{ShipOp2}|\boldsymbol{\alpha^A}, \boldsymbol{\beta^A}, ip, b_0^{\mathcal{M}}\right)\Big)\mathcal{P}_A\left(\boldsymbol{\alpha^A}, \boldsymbol{\beta^A}, ip, b_0^{\mathcal{M}}\right) \\
&\qquad + Pr(\text{Case B}) \sum_{(\boldsymbol{\alpha^B}, \boldsymbol{\beta^B}, ip, b_0^{\mathcal{M}}) \in \Theta_B} \Big(\mathcal{P}_B\Big(\text{ShipOp1}|\boldsymbol{\alpha^B}, b_0^{\mathcal{M}}\Big) \\
&\qquad + \mathcal{P}_B\left(\text{ShipOp2}|\boldsymbol{\alpha^B}, \boldsymbol{\beta^B}, ip, b_0^{\mathcal{M}}\right)\Big)\mathcal{P}_B\left(\boldsymbol{\alpha^B}, \boldsymbol{\beta^B}, ip, b_0^{\mathcal{M}}\right), \quad (3.16)
\end{aligned}$$

where Θ_A and Θ_B are the sets of feasible combinations of values in the support of the joint PMF for Case A and Case B, respectively. Note that the possible values of the realizations of the random variables are given by $\boldsymbol{\alpha^i} \geq 0$, $\boldsymbol{\beta^i} \geq 0$, $R_0 + 1 \leq ip \leq R_0 + Q_0$,

$b_0^{\mathcal{M}} \geq 0$. $\boldsymbol{\alpha^i}$, $\boldsymbol{\beta^i}$, and ip are random vectors/variables independent of each other while $b_0^{\mathcal{M}}$ depends on the other random variables, as shown in Equation (3.25).

In Section 3.3.2, we derive the conditional probabilities of ShipOp1 and ShipOp2 given demand realizations in the different intervals $(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i})$, the warehouse inventory position (ip), and the number of backorders at the warehouse for retailer group $\mathcal{M}$ $(b_0^{\mathcal{M}})$. The random variable T_I is accounted for when deriving the joint PMF of these realizations, $\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}})$, for Case A and Case B, respectively. This is done in Section 3.3.3.

3.3.2 Deriving the Conditional Probability for Shipment Option 1 and Shipment Option 2

In this section, we derive the conditional probability that the warehouse dispatches the considered n^{th} unit on the time-based shipment at $\tau' - T_I$ or before, referred to as $\mathcal{P}_i(\text{ShipOp1}|\boldsymbol{\alpha^i}, b_0^{\mathcal{M}})$, and the conditional probability that the warehouse dispatches the considered n^{th} unit after $\tau' - T_I$ but before τ' on a potential quantity-based shipment, referred to as $\mathcal{P}_i(\text{ShipOp2}|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}})$. The analysis is performed conditioned on the realizations: $\boldsymbol{\alpha^i}$, $\boldsymbol{\beta^i}$, ip, and $b_0^{\mathcal{M}}$ for $i \in \{\text{A,B}\}$.

The Probability $\mathcal{P}_i(\text{ShipOp1}|\boldsymbol{\alpha^i}, b_0^{\mathcal{M}})$:

We first focus on $\mathcal{P}_i(\text{ShipOp1}|\boldsymbol{\alpha^i}, b_0^{\mathcal{M}})$, i.e., the probability that the considered n^{th} unit is qualified for shipment prior to $\tau' - T_I$. As mentioned, all units qualified for shipment at or before the dispatch of a time-based shipment will be on this shipment or an earlier shipment. Consequently, the event ShipOp1 occurs if the considered n^{th} unit is both demanded by a retailer and available at the warehouse no later than $\tau' - T_I$, i.e., the time of the most recent time-based shipment before τ'.

We start by considering whether the considered n^{th} unit is demanded before or after $\tau' - T_I$. Recall that the considered n^{th} unit is the n^{th} most recent unit ordered by retailer group $\mathcal{M}$ prior to τ'. From Figure 3.4 and Figure 3.5, we can see that this considered n^{th} unit is demanded before $\tau' - T_I$ if and only if $n > \alpha_1^A + \alpha_2^A$ in Case A and $n > \alpha_1^B$ in Case B.

Next, we consider if the considered n^{th} unit is available prior to time $\tau' - T_I$. At time $\tau' - T_I$, there are $b_0^{\mathcal{M}}$ backordered units at the warehouse for retailer group $\mathcal{M}$. These backordered units will satisfy the $b_0^{\mathcal{M}}$ demands with the lowest priority among the demands occurring before $\tau' - T_I$ (i.e., the most recent demands before $\tau' - T_I$).

The considered n^{th} unit must be available at the warehouse at $\tau' - T_I$ if it is not backordered at this time. This must be the case if the considered n^{th} unit is ordered before the $b_0^{\mathcal{M}}$ backordered units. The conditions under which this is true differ between Case A and Case B. In Case A, the condition is true if and only if $n > \alpha_1^A + \alpha_2^A + b_0^{\mathcal{M}}$, whereas in Case B, the condition is $n > \alpha_1^B + b_0^{\mathcal{M}}$. When interpreting these conditions, it is essential to keep in mind that n denotes the priority of the considered n^{th} unit, counting backward in time from τ', $b_0^{\mathcal{M}}$ correspond to the most recent demands from retailer group $\mathcal{M}$ prior to $\tau' - T_I$, and the demand at retailer group $\mathcal{M}$ in time interval $(\tau' - T_I, \tau']$ is $\alpha_1^A + \alpha_2^A$ for Case A and α_1^B for Case B. Note that by definition, $b_0^{\mathcal{M}} \geq 0$. We summarize these arguments in the following Lemma.

Lemma 1. *Assume that the realizations $b_0^{\mathcal{M}}$, $\boldsymbol{\alpha^A}$, and $\boldsymbol{\alpha^B}$ are known for Case A and Case B, respectively; then the probability that the considered n^{th} unit is dispatched with or before the time-based shipment at time $\tau' - T_I$ is*

$$Pr\left(\textit{ShipOp1}|\boldsymbol{\alpha^A}, b_0^{\mathcal{M}}\right) = \begin{cases} 1 & \textit{if } n > \alpha_1^A + \alpha_2^A + b_0^{\mathcal{M}} \\ 0 & \textit{otherwise} \end{cases} \qquad \textit{for Case A} \tag{3.17}$$

$$Pr\left(\textit{ShipOp1}|\boldsymbol{\alpha^B}, b_0^{\mathcal{M}}\right) = \begin{cases} 1 & \textit{if } n > \alpha_1^B + b_0^{\mathcal{M}} \\ 0 & \textit{otherwise.} \end{cases} \qquad \textit{for Case B.} \tag{3.18}$$

The Probability $\mathcal{P}_i$(ShipOp2$|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}}$):

In this section, we focus on $\mathcal{P}_i(\text{ShipOp2}|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}})$, i.e., the probability that the warehouse dispatches the considered n^{th} unit on a quantity-based shipment after $\tau' - T_I$ and before τ'. As noted before, this requires that the last unit in a quantity-based shipment, including the considered n^{th} unit, is qualified for shipment before time τ'.

We first consider if the last unit is demanded by retailer group $\mathcal{M}$ before τ'. Recall the priority list of units ordered before τ' by retailer group $\mathcal{M}$, where the unit that will satisfy the most recent demand before τ' has priority 1, the unit that will satisfy the second most recent demand before τ' has priority 2, and so on. We define the priority of the last unit as q. Note that $q > 0$ means that the demand that this unit will satisfy takes place before τ'. Conversely, if $q \leq 0$, the demand for the last unit occurs after τ', which means it can never be dispatched in time for the event ShipOp2 to occur. Also note that if the last unit in a shipment with $Q^{\mathcal{M}}$ units has priority q, the considered n^{th} unit belongs to the same quantity-based shipment if and only if $n - q \leq Q^{\mathcal{M}} - 1$. For example, if the last unit in a shipment with $Q^{\mathcal{M}} = 4$ units has priority $q = 5$ the other units in the batch will have priorities 6,7 and 8. Thus, in order for the considered n^{th} unit to be part of this quantity-based shipment, n must equal 5,6,7 or 8. A general expression for q given n is provided in Proposition 2. A proof of Proposition 2 is provided in Appendix 3.7.3.

Proposition 2. *Given that the considered n^{th} unit was not dispatched before or at $\tau' - T_I$, the last unit in a possible quantity-based shipment, including the considered n^{th} unit, corresponds to priority q for retailer group $\mathcal{M}$ at time τ', where*

$$q = \begin{cases} n - Q^{\mathcal{M}} + 1 + \mathrm{mod}_{Q^{\mathcal{M}}}\left(b_0^{\mathcal{M}} + \alpha_1^A + \alpha_2^A - n\right) & \text{for Case A} \\ n - Q^{\mathcal{M}} + 1 + \mathrm{mod}_{Q^{\mathcal{M}}}\left(b_0^{\mathcal{M}} + \alpha_1^B - n\right) & \text{for Case B.} \end{cases} \tag{3.19}$$

It remains to be determined if the last unit is available for shipment from the warehouse before or at τ'. This will be the case if and only if the warehouse orders the last unit, with priority $q > 0$, no later than $\tau' - L_0$. As the warehouse faces demand from all retailer groups, the time of the warehouse order depends on the total demand at the warehouse (i.e., the demand sequence from all retailers). It also depends on the number of backorders at time $\tau' - T_I$, $b_0^{\mathcal{M}}$, since this is influenced by the conditioned warehouse inventory position at time $\tau' - T_I - L_0$, $IP_0(\tau' - T_I - L_0) = ip$. To account for these dependencies, we determine when the last unit is ordered by the warehouse relative to the time instance $\tau' - T_I - L_0$.

To proceed, we start by determining which unit the last unit corresponds to, counting the demands of all retailer groups, starting from $\tau' - T_I - L_0$ until (and including) the last unit. We define

$\Psi(q)$ The number of system demands that occur from $\tau' - T_I - L_0$ until (and including) the demand for the last unit with priority q takes place at retailer group $\mathcal{M}$. It should be noted that $\Psi(q) < 0$ if the last unit is demanded before $\tau' - T_I - L_0$, i.e., $q > \alpha_1^i + \alpha_2^i + \alpha_3^i$.

In other words, $\Psi(q)$ identifies the last unit in terms of the system demands that take place between $\tau' - T_I - L_0$ and the retailer order from retailer group $\mathcal{M}$ that demands the last unit. For a given warehouse inventory position at time $\tau' - T_I - L_0$, $\Psi(q)$ makes it possible to relate the retailer demand of the last unit to the warehouse order containing this unit. By defining the time of the retailer order of the last unit as t_{ro} ($ro =$ retailer order), it follows that for $\Psi(q) > 0$, $\Psi(q)$ denotes the system demand from $\tau' - T_I - L_0$ up until (and including) the demand that occurs at t_{ro}, denoted by $D_0(\tau' - T_I - L_0, t_{ro})$. If $\Psi(q) < 0$, we have $\Psi(q) = -\tilde{D}_0(t_{ro}, \tau' - T_I - L_0)$, where $\tilde{D}(t_{ro}, \tau' - T_I - L_0)$ denotes the system demand from (and including) the demand that occurs at t_{ro} to $\tau' - T_I - L_0$. Figure 3.6 visualizes $\Psi(q) > 0$ in a timeline including the important time points.

However, since we do not know the sequence of demands from the respective retailer groups within an interval, we do not know how many of the demands from the other retailer groups occur before (and after) the demand of the last unit. As a consequence, $\Psi(q)$ is a stochastic variable. We derive the distribution of $\Psi(q)$ in Proposition 3 and Proposition 4.

For given realizations of $\Psi(q) = \psi$ and $IP_0(\tau' - T_I - L_0) = ip$, we can determine the time when the warehouse orders the last unit, denoted by t_{wo}, and relate that to time t_{ro} when it is ordered by a retailer in group $\mathcal{M}$. This is done by determining how many system demands occur after time t_{wo} (when the warehouse orders the unit satisfying the ψ^{th} system demand after $\tau' - T_I - L_0$) until time t_{ro} (when the ψ^{th} system demand is ordered by a retailer in group $\mathcal{M}$ counting from $\tau' - T_I - L_0$). We define this quantity by $\phi(\psi|ip)$.

$\phi(\psi|ip)$ The number of system demands that occur after the warehouse orders the unit that will satisfy the ψ^{th} system demand after $\tau' - T_I - L_0$, until (and including) the ψ^{th} system demand, given that $IP_0(\tau' - T_I - L_0) = ip$. It should be noted that $\phi(\psi|ip) \leq 0$ if the warehouse orders the replenishment batch, including the last unit, after (or at the same time) the retailer group $\mathcal{M}$ orders the last unit.

$\phi(\psi|ip)$ may take values in $\{R_0 + 1, R_0 + 2, \ldots, R_0 + Q_0\}$. Thus, if the warehouse uses a negative reorder point, $\phi(\psi|ip)$ may take both positive and negative values. A positive value implies that the warehouse orders the last unit from the outside supplier before the retailer orders it from the warehouse. Conversely, a negative value indicates that the warehouse orders the last unit after the retailer orders it. Similarly, if $\phi(\psi|ip) = 0$, the

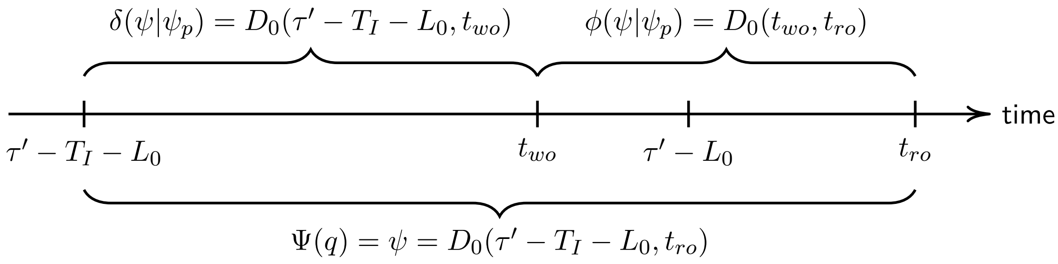

Figure 3.6: A timeline illustrating how $\delta(\psi|ip), \Psi(q) = \psi$ and $\phi(\Psi(q)|ip)$ relate to each other when all variables are positive.

warehouse order and the retailer order of the last unit occur simultaneously. Consequently, if $R_0 \geq -1$, the retailer will always order units from the warehouse after or at the same time as the warehouse orders the same units from the outside supplier.

Lastly, we will determine if t_{wo} takes place before or after $\tau' - L_0$, meaning that the last unit will be available at the warehouse at time τ'. To do so, we define the time of the warehouse order in terms of the number of demands at the warehouse after $\tau' - T_I - L_0$ up until (and including) the demand that triggers the warehouse order. This quantity is referred to as

$\delta(\psi|ip)$ The number of system demands between $\tau' - T_I - L_0$, until (and including) the demand that triggers the warehouse order that includes the last unit, given that $IP_0(\tau' - T_I - L_0) = ip$. It should be noted that $\delta(\psi|ip) < 0$ if the warehouse orders the replenishment batch, including the last unit, before $\tau' - T_I - L_0$.

Thus, if $\delta(\psi|ip) > 0$, $\delta(\psi|ip) = D_0(\tau' - T_I - L_0, t_{wo})$. If $\delta(\psi|ip) < 0$, we have $\delta(\psi|ip) = -\tilde{D}_0(t_{wo}, \tau' - T_I - L_0)$. By definition of $\Psi(q)$ and $\phi(\psi|ip)$, we have

$$\delta(\Psi(q)|ip) = \Psi(q) - \phi(\Psi(q)|ip). \tag{3.20}$$

Figure 3.6 illustrates a timeline for how $\delta(\psi|ip), \Psi(q) = \psi$ and $\phi(\Psi(q)|ip)$ relate to each other if all variables are positive.

Since the number of demands at the warehouse in the interval $(\tau' - T_I - L_0, \tau' - L_0]$ is known from $\boldsymbol{\alpha^i}$ and $\boldsymbol{\beta^i}$, we can compare this quantity to $\delta(\psi|ip)$ to determine if the warehouse order of the last unit at t_{wo} takes place before or after $\tau' - L_0$. From Figure 3.4, we can see that, for Case A $t_{wo} \leq \tau' - L_0$ if and only if $\delta(\psi|ip) \leq \alpha_3^A + \beta_3^A + \alpha_2^A + \beta_2^A$. Similarly, for Case B, Figure 3.5 illustrates that $t_{wo} \leq \tau' - L_0$ if and only if $\delta(\psi|ip) \leq \alpha_3^B + \beta_3^B$. Moreover, since ShipOp2 can never occur if ShipOp1 occurs, ShipOp2 is only relevant if $n \leq \alpha_1^A + \alpha_2^A + b_0^{\mathcal{M}}$ in Case A and $n \leq \alpha_1^A + b_0^{\mathcal{M}}$ in Case B. These arguments are summarized in Lemma 2.

Lemma 2. *Assume that $b_0^{\mathcal{M}}$, ip, $\boldsymbol{\alpha^i}$, and $\boldsymbol{\beta^i}$ are known for Case A and Case B, respectively; then the probability that the considered n^{th} unit is dispatched with a quantity-based shipment after $\tau' - T_I$ and before τ' is given for Case A by Equation (3.21) and Case B by Equation (3.22).*

$$Pr\Big(\textit{ShipOp2}|\boldsymbol{\alpha^A}, \boldsymbol{\beta^A}, ip, b_0^{\mathcal{M}}\Big) = \begin{cases} Pr\left(\delta(\Psi(q)|ip) \leq \alpha_3^A + \alpha_2^A + \beta_3^A + \beta_2^A\right) & \textit{if } q > 0 \textit{ and } n \leq \alpha_1^A + \alpha_2^A + b_0^{\mathcal{M}} \\ 0 & \textit{otherwise} \end{cases} \tag{3.21}$$

$$Pr\Big(\textit{ShipOp2}|\boldsymbol{\alpha^B}, \boldsymbol{\beta^B}, ip, b_0^{\mathcal{M}}\Big) = \begin{cases} Pr\left(\delta(\Psi(q)|ip) \leq \alpha_3^B + \beta_3^B\right) & \textit{if } q > 0 \textit{ and } n \leq \alpha_1^B + b_0^{\mathcal{M}} \\ 0 & \textit{otherwise} \end{cases} \tag{3.22}$$

The remaining part of this section explains how to determine $\phi(\psi|ip)$ and $Pr(\Psi(q) = \psi)$ using Proposition 3, Proposition 4, and Proposition 5. The detailed derivations of the last two propositions are deferred to Appendix 3.7.4 and Appendix 3.7.5.

We begin with $\Psi(q)$, i.e., the total number of demands in the system from $\tau' - T_I - L_0$ up until (and including) the demand of the last unit with priority q. The expression for this quantity depends on the time interval in which the last unit is ordered by retailer group $\mathcal{M}$ (see Figure 3.4 and Figure 3.5).

Proposition 3. *The number of system demands from $\tau' - T_I - L_0$ until (and including) the demand for the last unit with priority q was ordered by retailer group $\mathcal{M}$ is given as*

$$\Psi(q) = \begin{cases} \alpha_3^i + \alpha_2^i + \alpha_1^i - q + 1 + \beta_3^i + \beta_2^i + X_{\beta_1^i} & \text{if } 0 < q \le \alpha_1^i \\ \alpha_3^i + \alpha_2^i + \alpha_1^i - q + 1 + \beta_3^i + X_{\beta_2^i} & \text{if } \alpha_1^i < q \le \alpha_1^i + \alpha_2^i \\ \alpha_3^i + \alpha_2^i + \alpha_1^i - q + 1 + X_{\beta_3^i} & \text{if } \alpha_1^i + \alpha_2^i < q \le \alpha_1^i + \alpha_2^i + \alpha_3^i \\ \alpha_3^i + \alpha_2^i + \alpha_1^i - q - X_{\beta_4^i} & \text{if } q > \alpha_1^i + \alpha_2^i + \alpha_3^i, \end{cases} \tag{3.23}$$

where

$X_{\beta_j^i}$ *for $j \in \{1, 2, 3\}$ is the (random) number of demands among the β_j^i demands from the other retailer groups $\overline{\mathcal{M}}$ in the same time interval as the last unit was ordered, that occur before retailer group $\mathcal{M}$ orders the last unit. Note that these random variables are only relevant for $q \le \alpha_1^i + \alpha_2^i + \alpha_3^i$,*

$X_{\beta_4^i}$ *is the (random) number of demands from the other retailer groups $\overline{\mathcal{M}}$ that have been ordered after the last unit but before $\tau' - T_I - L_0$. Note that these random variables are only relevant for $q > \alpha_1^i + \alpha_2^i + \alpha_3^i$.*

The fourth expression in Equation (3.23) considers the instance when the last unit is demanded prior to $\tau' - T_I - L_0$, i.e., before the three time intervals. Recall that $\Psi(q)$ takes negative values in this case. Moreover, note that $\Psi(q)$ is stochastic, as the expression for $\Psi(q)$ involves the random variable $X_{\beta_j^i}$. The distribution of $X_{\beta_j^i}$ is presented in Proposition 4.

Proposition 4. *The random variable $X_{\beta_j^i}$ can be obtained as*

$$X_{\beta_j^i} = \xi + X, \tag{3.24}$$

where ξ is a constant and $X \sim Betabin(k, \alpha, \beta)$, and $Betabin$ denotes the beta-binomial distribution. The value of ξ and the parameters k, α, and β differ depending on different cases with respect to ip and α_j^i and β_j^i. The parameters for every special case can be found in Table 3.2.

In Table 3.2, we consider different cases depending on how many backordered demands there are at time $\tau' - T_I$, i.e., $B_0(\tau' - T_I) = b_0$. Note that the realization b_0 is uniquely

Table 3.2: The table provides parameters (ξ, k, α, and β) for the distribution of $X_i^j = \xi + X$, where $X \sim Betabin(k, \alpha, \beta)$ for Case A and Case B, respectively. Note that in case $\xi = 0$, $X = X_i^j \sim Betabin(k, \alpha, \beta)$.

			$X_{\beta_1^i} = \xi + X$	$X_{\beta_2^i} = \xi + X$	$X_{\beta_3^i} = \xi + X$	$X_{\beta_4^i} = \xi + X$
A	$b_0 \leq \alpha_3^A + \beta_3^A$	ξ	0	0	$\beta_3^B - (b_0 - b_0^{\mathcal{M}})$	-
		k	$\alpha_1^A - q + 1$	$\alpha_1^A + \alpha_2^A - q + 1$	$\alpha_1^A + \alpha_2^A + b_0^{\mathcal{M}} - q + 1$	-
		α	α_1^A	α_2^A	$b_0^{\mathcal{M}}$	-
		β	β_1^A	β_2^A	$b_0 - b_0^{\mathcal{M}}$	-
	$b_0 > \alpha_3^A + \beta_3^A$	ξ	0	0	0	0
		k	$\alpha_1^A - q + 1$	$\alpha_1^A + \alpha_2^A - q + 1$	$\alpha_1^A + \alpha_2^A + \alpha_3^A - q + 1$	$q - \alpha_1^A - \alpha_2^A - \alpha_3^A$
		α	α_1^A	α_2^A	α_3^A	$b_0^{\mathcal{M}} - \alpha_3^A$
		β	β_1^A	β_2^A	β_3^A	$(b_0 - b_0^{\mathcal{M}}) - \beta_3^A$
B	$b_0 \leq \alpha_2^B + \beta_2^B$	ξ	0	$\beta_2^B - (b_0 - b_0^{\mathcal{M}})$	-	-
		k	$\alpha_1^B - q + 1$	$\alpha_1^B + b_0^{\mathcal{M}} - q + 1$	-	-
		α	α_1^B	$b_0^{\mathcal{M}}$	-	-
		β	β_1^B	$b_0 - b_0^{\mathcal{M}}$	-	-
	$\alpha_2^B + \beta_2^B < b_0 \leq \alpha_2^B + \beta_2^B + \alpha_3^B + \beta_3^B$	ξ	0	0	$\beta_2^B + \beta_3^B - (b_0 - b_0^{\mathcal{M}})$	-
		k	$\alpha_1^B - q + 1$	$\alpha_1^B + \alpha_2^B - q + 1$	$\alpha_1^B + \alpha_2^B + b_0^{\mathcal{M}} - q + 1$	-
		α	α_1^B	α_2^B	$b_0^{\mathcal{M}} - \alpha_2^B$	-
		β	β_1^B	β_2^B	$(b_0 - b_0^{\mathcal{M}}) - \beta_2^B$	-
	$b_0 > \alpha_2^B + \beta_2^B + \alpha_3^B + \beta_3^B$	ξ	0	0	0	0
		k	$\alpha_1^A - q + 1$	$\alpha_1^B + \alpha_2^B - q + 1$	$\alpha_1^B + \alpha_2^B + \alpha_3^B - q + 1$	$q - \alpha_1^B - \alpha_2^B - \alpha_3^B$
		α	α_1^B	α_2^B	α_3^B	$b_0^{\mathcal{M}} - \alpha_3^B - \alpha_2^B$
		β	β_1^B	β_2^B	β_3^B	$(b_0 - b_0^{\mathcal{M}}) - \beta_3^A - \beta_2^B$

given from the realization of $IP_0(\tau' - T_I - L_0) = ip$ and the demand realizations $\boldsymbol{\alpha^i}$ and $\boldsymbol{\beta^i}$, especially the demand realizations during the supply lead time from the outside supplier to the warehouse. Since we have conditioned the analysis on fixed values of ip, $\boldsymbol{\alpha^i}$, and $\boldsymbol{\beta^i}$, we can formulate

$$b_0 = \begin{cases} \left(ip - \alpha_3^A - \beta_3^A\right)^- & \text{for Case A} \\ \left(ip - \alpha_2^B - \alpha_3^B - \beta_2^B - \beta_3^B\right)^- & \text{for Case B.} \end{cases} \tag{3.25}$$

Finally, Proposition 5 shows how to obtain $\phi(\psi|ip)$, i.e., the number of system demands that occur between the time when the warehouse orders the unit that will satisfy the ψ^{th} system demand at retailer group $\mathcal{M}$ after $\tau' - T_I - L_0$ and the time when it is demanded by retailer group $\mathcal{M}$.

Proposition 5. *Consider an arbitrary unit ordered by the warehouse at some time t_{wo}. The retailer demand for this unit occurs at time t_{ro} and corresponds to the ψ^{th} system demand after time $\tau' - T_I - L_0$. Assume that the inventory position at the warehouse at time $\tau' - T_I - L_0$ is equal to ip. Then the demand that triggered the warehouse order at t_{wo} occurred $\phi(\psi|ip)$ system demands prior to the demand at time t_{ro}.*

$$\phi(\psi|ip) = R_0 + Q_0 - \mathrm{mod}_{Q_0}(ip - \psi) \tag{3.26}$$

3.3.3 The Joint Probability Mass Function of $\boldsymbol{\alpha}, \boldsymbol{\beta}, ip$, and $b_0^{\mathcal{M}}$

This section explains the approach for determining the joint PMF of the random variables with the realizations $\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip$, and $b_0^{\mathcal{M}}$, conditioned on Case $i \in$ {A,B}, i.e., $\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}})$.

A first observation is that $IP_0(\tau' - T_I - L_0)$ is independent of the demand that occurs after time $\tau' - T_I - L_0$ in the three time intervals we consider, see Figure 3.4 and Figure 3.5. Thus, we first derive the PMF of the demand and the inventory position individually. Subsequently, we derive the marginal PMF of $B_0^{\mathcal{M}}(\tau' - T_I)$ given $\boldsymbol{\alpha^i}$, $\boldsymbol{\beta^i}$, and ip. Finally, we obtain the joint PMF by combining these results:

$$\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip, b_0^{\mathcal{M}}) = \mathcal{P}_i(b_0^{\mathcal{M}}|\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}, ip)\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i})\mathcal{P}_i(ip). \tag{3.27}$$

Let us first consider the PMF of the inventory position at the warehouse, $\mathcal{P}_i(ip)$. As previously mentioned, ip represents a realization of $IP_0(\tau' - T_I - L_0)$. Because the replenishment orders in the system are independent of the shipment consolidation policy, the inventory position at the warehouse behaves in exactly the same way as in a system without shipment consolidation. It has been well established in the literature (see e.g., Axsäter (2015)) that the inventory position in such a system is uniformly distributed on the integers $[R_0+1, R_0+2, \ldots, R_0+Q_0]$ not only for Poisson demand. Hence,

$$\mathcal{P}_i(ip) = \frac{1}{Q_0} \qquad \forall\ ip \in \{R_0+1, R_0+2, \ldots, R_0+Q_0\}. \tag{3.28}$$

We now turn to the demand in the three time intervals, $\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i})$. If the lengths of these intervals are fixed, the number of demands in the different intervals is independent of each other. However, when the lengths of the intervals are stochastic, there is a dependency between α_j^i and β_j^i through their mutual dependency of the random interval length. To facilitate the derivation of these random variables, we first focus on the system demand in Interval j, denoted by $\gamma_j^i = \alpha_j^i + \beta_j^i$. If we derive the joint distribution $\mathcal{P}_i(\boldsymbol{\gamma^i}) = \mathcal{P}_i(\gamma_1^i, \gamma_2^i, \gamma_3^i)$, the number of demands among the γ_j^i demands that belong to retailer group $\mathcal{M}$ can be obtained by binomial disaggregation:

$$\mathcal{P}_i(\alpha_j^i|\gamma_j^i) = \binom{\gamma_j^i}{\alpha_j^i}\left(\frac{\lambda^{\mathcal{M}}}{\Lambda}\right)^{\alpha_j^i}\left(1-\frac{\lambda^{\mathcal{M}}}{\Lambda}\right)^{\gamma_j^i-\alpha_j^i} \tag{3.29}$$

β_j^i is then obtained by $\beta_j^i = \gamma_j^i - \alpha_j^i$. The relationships between α_j^i, β_j^i, and γ_j^i allows us to formulate

$$\mathcal{P}_i(\boldsymbol{\alpha^i}, \boldsymbol{\beta^i}) = \mathcal{P}_i(\alpha_1^i|\gamma_1^i)\mathcal{P}_i(\alpha_2^i|\gamma_2^i)\mathcal{P}_i(\alpha_3^i|\gamma_3^i)\mathcal{P}_i(\gamma_1^i, \gamma_2^i, \gamma_3^i). \tag{3.30}$$

We now consider the joint PMF of the system demands in each of the three intervals, i.e., $\mathcal{P}_i(\gamma_1^i, \gamma_2^i, \gamma_3^i)$. This derivation differs depending on which of the two cases (Case A and Case B) we consider. In Case A, we have three disjoint time intervals. Two of them (Interval 1 and Interval 3) are of fixed length L_0, and the third (Interval 2) is of stochastic length $T_I - L_0$, see Figure 3.4. The fixed interval length of all but one of the intervals means that γ_1^A, γ_2^A, and γ_3^A are realizations from independent random variables.

In Case B, we can see from Figure 3.5 that we have two intervals of length T_I (Interval 1 and Interval 3) and a third interval of length $L_0 - T_I$. Since the lengths of all three intervals depend on T_I, the number of demands in these intervals (γ_1^B, γ_2^B, and γ_3^B) are all dependent through their mutual dependency on T_I.

Let us start with the simpler Case A. As mentioned above, there are two time intervals of fixed length L_0. Since the demand occurs according to Poisson processes, the total number of demands in the system during the fixed time intervals (Interval 1 and Interval 3) are independent and identically distributed with a Poisson distribution with rate ΛL_0.

$$\mathcal{P}_A(\gamma_1^A, \gamma_3^A) = \mathcal{P}_A(\gamma_1^A)\mathcal{P}_A(\gamma_3^A) = \frac{\mathrm{e}^{-2\Lambda L_0}(\Lambda L_0)^{(\gamma_1^A+\gamma_3^A)}}{(\gamma_1^A+\gamma_3^A)!} \tag{3.31}$$

Moreover, the number of demands in these time intervals is independent of the number of demands in Interval 2. In other words, $\mathcal{P}_A(\gamma_1^A, \gamma_2^A, \gamma_3^A) = \mathcal{P}_A(\gamma_1^A, \gamma_3^A)\mathcal{P}_A(\gamma_2^A)$. To obtain $\mathcal{P}_A(\gamma_2^A)$, i.e., $Pr(D^A(\tau'-T_I, \tau'-L_0) = \gamma_2^A)$ we first condition on the value of T_I and then take the expectation across all possible values, i.e., for group $\mathcal{M}$, $T_I \in \left(L_0, T^{\mathcal{M}}\right)$. Note that we only consider Case A, i.e., $T_I > L_0$. We formulate the following Proposition 6, with its proof provided in Appendix 3.7.6.

Proposition 6. *For Case A, the PMF $Pr(D^A(\tau' - T_I, \tau' - L_0) = \gamma_2^A) = \mathcal{P}_A(\gamma_2^A)$ is given by*

$$\mathcal{P}_A(\gamma_2^A) = \frac{1}{\Lambda(T^{\mathcal{M}} - L_0)}\left[1 - F\left(\gamma_2^A, \Lambda(T^{\mathcal{M}} - L_0)\right)\right], \tag{3.32}$$

where $F^{\lambda}(k)$ is the cumulative distribution function of the Poisson distribution with parameter λ, i.e.,

$$F^{\lambda}(k) = \mathrm{e}^{-\lambda}\sum_{j=0}^{k}\frac{\lambda^j}{j!}. \tag{3.33}$$

In Case B, the demand in all three time intervals are mutually dependent on T_I. We therefore directly formulate the joint PMF of γ_1^B, γ_2^B, and γ_3^B in Proposition 7. The proposition is derived in a similar way, conditioning on the random variable T_I and then taking the expectation across all $T_I \in [0, \min(L_0, T^m)]$. The derivation is found in Appendix 3.7.7.

Proposition 7. *For Case B, the joint PMF* $Pr(D^B(\tau'-T_I,\tau')=\gamma_1^B, D^B(\tau'-L_0,\tau'-T_I)=\gamma_2^B, D^B(\tau'-T_I-L_0,\tau'-L_0)=\gamma_3^B)=\mathcal{P}_B(\gamma_1^B,\gamma_2^B,\gamma_3^B)=\mathcal{P}_B(\boldsymbol{\gamma^B})$ *is given by*

$$\mathcal{P}_B(\gamma_1^B,\gamma_2^B,\gamma_3^B)=\frac{1}{b}\frac{(\Lambda L_0)^{\gamma_2^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\frac{e^{-\Lambda L_0}}{\Lambda} \cdot\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}\left(\frac{-1}{\Lambda L_0}\right)^k(\gamma_1^B+\gamma_3^B+k)!\left[1-F^{\Lambda b}(\gamma_1^B+\gamma_3^B+k)\right], \tag{3.34}$$

where $b=\min(L_0,T^{\mathcal{M}})$ *and* $F^{\lambda}(k)$ *is defined as in Equation (3.33).*

We finally turn to the number of backorders for retailer group $\mathcal{M}$ at time $\tau'-T_I$, i.e., $B_0^{\mathcal{M}}(\tau'-T_I)=b_0^{\mathcal{M}}$. Recall that $B_0(\tau'-T_I)=b_0$ is directly obtained from ip and the system demand during $(\tau'-T_I-L_0,\tau'-T_I)$ as shown in Equation (3.25).

The proportion among the b_0 backorders that are associated with retailer group $\mathcal{M}$ depends on the proportions between $\boldsymbol{\alpha^i}$ and $\boldsymbol{\beta^i}$. To handle this dependency, we formulate the Proposition 8, with its derivation provided in Appendix 3.7.8.

Proposition 8. *The PMF* $\mathcal{P}_i(b_0^{\mathcal{M}}|\boldsymbol{\alpha^i},\boldsymbol{\beta^i},ip)$ *is given by*

$$\mathcal{P}_i(b_0^{\mathcal{M}}|\boldsymbol{\alpha^i},\boldsymbol{\beta^i},ip)=Pr(X=b_0^{\mathcal{M}}-\kappa), \tag{3.35}$$

where κ *is a constant and* X *is a random variable. For Case* A*,* κ *and* X *is given by*

$$\kappa=\begin{cases}0 & \text{for } b_0\le\alpha_3^A+\beta_3^A\\ \alpha_3^A & \text{for } b_0>\alpha_3^A+\beta_3^A\end{cases} \tag{3.36}$$

and

$$X\sim\begin{cases}Hyp(\alpha_3^A+\beta_3^A,\alpha_3^A,b_0) & \text{for } b_0\le\alpha_3^A+\beta_3^A\\ Bin(b_0-\alpha_3^A-\beta_3^A,\frac{\lambda^{\mathcal{M}}}{\Lambda}) & \text{for } b_0>\alpha_3^A+\beta_3^A,\end{cases} \tag{3.37}$$

where Hyp *and* Bin *denote the hypergeometric and binomial distribution, respectively. For Case* B *we have*

$$\kappa=\begin{cases}0 & \text{for } b_0\le\alpha_2^B+\beta_2^B\\ \alpha_2^B & \text{for } \alpha_2^B+\beta_2^B<b_0\le\alpha_2^B+\beta_2^B+\alpha_3^B+\beta_3^B\\ \alpha_2^B+\alpha_3^B & \text{for } b_0>\alpha_2^B+\beta_2^B+\alpha_3^B+\beta_3^B\end{cases} \tag{3.38}$$

and

$$X \sim \begin{cases} Hyp(\alpha_2^B + \beta_2^B, \alpha_2^B, b_0) & \textit{for } b_0 \leq \alpha_2^B + \beta_2^B \\ Hyp(\alpha_3^B + \beta_3^B, \alpha_3^B, b_0 - \alpha_2^B - \beta_2^B) & \textit{for } \alpha_2^B + \beta_2^B < b_0 \leq \alpha_2^B + \beta_2^B + \alpha_3^B + \beta_3^B \\ Bin(b_0 - \alpha_2^B - \beta_2^B - \alpha_3^B - \beta_3^B, \frac{\lambda^{\mathcal{M}}}{\Lambda}) & \textit{for } b_0 > \alpha_2^B + \beta_2^B + \alpha_3^B + \beta_3^B. \end{cases} \tag{3.39}$$

3.4 Optimization

In this section, we explain how the expected total cost function, $TC_M(R_0, \mathbf{S}, \mathbf{T}, \mathbf{Q})$, can be minimized with respect to the warehouse reorder level R_0, the retailers' base-stock levels $\mathbf{S} = \{S_1, S_2, \ldots, S_{\mathcal{N}}\}$, the shipment intervals $\mathbf{T} = \{T^1, T^2, \ldots, T^{\mathcal{M}}\}$, and the shipment quantities $\mathbf{Q} = \{Q^1, Q^2, \ldots, Q^{\mathcal{M}}\}$. It is easy to show by example that $TC_M(R_0, \mathbf{S}, \mathbf{T}, \mathbf{Q})$ is not jointly convex in these decision variables. However, for given R_0, $\mathbf{T}$, and $\mathbf{Q}$, TC_M is separable and convex in the retailers' base-stock levels, $S_i \; \forall \; i \in \{1, 2, \ldots, \mathcal{N}\}$. This follows as $\mathbf{S}$ does not affect the warehouse demand, replenishment, and allocation processes in steady-state. Consequently, the warehouse delays are also unaffected and the optimal base-stock levels $\mathbf{S}^* = (S_1^*, S_2^*, \ldots, S_{\mathcal{N}}^*)$ can be obtained for a given $R_0, \mathbf{T}$, and $\mathbf{Q}$. Using this property, TC_M is optimized by a bounded search over $R_0, \mathbf{T}$, and $\mathbf{Q}$.

The search space is bounded by the sets of feasible shipment intervals $T^m \in \mathcal{T}^m$ and shipment quantities $Q^m \in \mathcal{Q}^m \; \forall \; m \in \{1, 2, \ldots, \mathcal{M}\}$, and all relevant values of $R_0 \in \{R_0^l, R_0^l + 1, \ldots, R_0^u\}$, where R_0^l and R_0^u denote a lower and upper bound on the optimal reorder level R_0^*. $R_0^l = -Q_0$ is a well-known lower bound (see, e.g., Axsäter (1998) and Marklund (2011)). An upper bound, R_0^u, can be found when the warehouse stock-out probability is small enough for further decreases to be of no consequence for the other decision variables (see, e.g., Axsäter (1990) and Stenius et al. (2016)). Thus,

$$R_0^u = min\{R_0 : Pr(D_0(0, L_0) > R_0) < \varepsilon\}, \tag{3.40}$$

where ε is a small positive number close to zero.

The optimal warehouse reorder level R_0^*, the optimal retailers' base-stock levels $\mathbf{S}^* = (S_1^*, S_2^*, \ldots, S_{\mathcal{N}}^*)$, the optimal shipment intervals $\mathbf{T}^* = \{T^{1*}, T^{2*}, \ldots, T^{\mathcal{M}*}\}$, and the optimal consolidation quantities $\mathbf{Q}^* = \{Q^{1*}, Q^{2*}, \ldots, Q^{\mathcal{M}*}\}$ can be found by minimizing the expected total cost over the bounded search space.

3.5 Numerical Study

The main purpose of the numerical study is to provide insights regarding the benefits and total cost performance of the hybrid consolidation policy compared to the time-based and quantity-based counterparts when jointly optimizing inventory and shipment decisions in OWMR inventory distribution systems.

The numerical study is based on a test series consisting of 24 different problems. The inventory system consists of a single warehouse and four retailers in two consolidation groups. Retailers 1 and 2 belong to retailer group 1, while retailers 3 and 4 belong to group 2. Our investigation encompasses various combinations of the following system parameters: The order quantity at the warehouse, backorder, and shipment costs, and the transportation times from the warehouse to the retailers. The order quantity at the warehouse is set at two levels, $Q_0 \in \{1, 10\}$. We define two levels each for the backorder costs rate, $\boldsymbol{p} = (p_1, p_2, p_3, p_4)$, and for the fixed shipment cost for quantity-based shipments, $\boldsymbol{\omega}_q = (\omega_q^1, \omega_q^2)$. In order to investigate varying relationships between the fixed costs associated with time-based and quantity-based shipments, $\boldsymbol{\omega}_t = (\omega_t^1, \omega_t^2)$ is studied at three different levels, each depending on the value of ω_q^m. Specifically, we consider $\boldsymbol{p} \in \{(10, 10, 10, 10), (50, 50, 50, 50)\}$, $\boldsymbol{\omega}_q \in \{(5, 5), (50, 50)\}$, and $\boldsymbol{\omega}_t \in \{0.4\boldsymbol{\omega}_q, 0.7\boldsymbol{\omega}_q, \boldsymbol{\omega}_q\}$.

The unit holding costs rate, h_i, the demand rates, $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$, the supply lead time, L_0, and the transportation times from the warehouse to the retailers, $\boldsymbol{L} = (L_1, L_2, L_3, L_4)$, are set as follows: $h_i = 1 \; \forall \; i = \{0, 1, 2, 3, 4\}$, $\boldsymbol{\lambda} = (0.4, 0.6, 0.1, 0.9)$, $L_0 = 2$, and $\boldsymbol{L} = (1, 1.2, 1, 1.2)$.These values are considered as constants in our analysis.

Further, we define the sets for the possible shipment policy parameters as $\mathcal{T}^m = \{1, 2, 3, 5, 7, 10, \infty\}$ and $\mathcal{Q}^m = \{1, 3, 5, 10, 20, \infty\} \; \forall \; m \in \{1, 2\}$.

Our method allows the calculation of many hybrid policies (T^m, Q^m), but the pure quantity-based policy with $T^m \to \infty$ cannot be evaluated. Therefore, the results for pure quantity-based shipment policies are obtained based on the method presented in Malmberg and Marklund (2023).

Table 3.3: Relative increase in expected total cost of using a pure time-based or a pure quantity-based consolidation policy instead of the hybrid policy

		Pure time-based policy			Pure quantity-based policy		
Parameters	Value	Minimum cost increase in %	Maximum cost increase in %	Average cost increase in %	Minimum cost increase in %	Maximum cost increase in %	Average cost increase in %
Q_0	1	0.00	18.16	6.29	0.00	28.14	5.31
	10	0.00	14.81	6.65	0.00	26.56	6.70
$\boldsymbol{p}$	(10,10,10,10)	0.00	18.16	6.21	0.00	28.14	6.93
	(50,50,50,50)	0.37	17.22	6.73	0.00	21.20	5.07
$\boldsymbol{\omega}_q$	(5,5)	0.13	18.16	9.41	0.00	11.76	2.75
	(50,50)	0.00	12.48	3.53	0.00	28.14	9.25
$\boldsymbol{\omega}_t$	$0.4\boldsymbol{\omega}_q$	0.00	7.11	1.78	3.64	28.14	15.62
	$0.7\boldsymbol{\omega}_q$	0.00	9.47	4.56	0.00	4.65	2.39
	$\boldsymbol{\omega}_q$	7.92	18.16	13.07	0.00	0.00	0.00
Total		0.00	18.16	6.47	0.00	28.14	6.00

When analyzing the results from our numerical study, we first the cost performance of the proposed hybrid policy to pure time-based and quantity-based shipment consolidation policies. Table 3.3 summarizes the minimum, maximum, and average increase in expected total cost for the optimized OWMR system when using a pure time-based or a pure quantity-based consolidation policy instead of the dominanting hybrid policy. We can see that the performance gap between the hybrid policy and a pure time-based policy ranges from 0 % to 18.16 % with an average of 6.47 % in our study. Thus, the expected shipment, inventory holding and backorder costs for the optimized OWMR system is on average 6.47 % higher when using a pure time-based policy instead of the hybrid policy. Similarly, for the optimal quantity-based policy, the relative increase in expected total costs range from 0 % to 28.14 % with an average of 6.00 %. This shows that there are substantial cost benefits of using the hybrid policy.

Looking more carefully at the results for different fixed costs per shipment (i.e., $\boldsymbol{\omega}_t \in \{0.4\boldsymbol{\omega}_q, 0.7\boldsymbol{\omega}_q, \boldsymbol{\omega}_q\}$) in Table 3.3, we can see that if the fixed costs per shipment is the same for the time-based and quantity-based shipments, i.e., $\boldsymbol{\omega}_t = \boldsymbol{\omega}_q$, the optimal pure quantity-based policy performs very well with 0 % cost increase across all examples. However, in practice, using periodic dispatches often offer advantages from a coordination perspective both for the in-house material handling activities, and by providing opportunities to coordinate activities across multiple items and warehouses (e.g., by cross-docking or similar set-ups). This suggests that time-based shipments may have lower fixed costs than the quantity-based shipments, i.e., $\boldsymbol{\omega}_t < \boldsymbol{\omega}_q$, motivating the inclusions of such scenarios in the study. Table 3.3 shows that in problems where the

Table 3.4: Optimal decision variables for all 24 instances

Instance	Q_0	$\boldsymbol{p}_j$	$\boldsymbol{\omega}_q$	$\boldsymbol{\omega}_t$	R_0^*	T^{1*}	T^{2*}	Q^{1*}	Q^{2*}	S_1^*	S_2^*	S_3^*	S_4^*	TC
1	1	(10,10,10,10)	(5,5)	$0.4\boldsymbol{\omega}_q$	3	1	2	10	5	2	3	1	5	14.10
2				$0.7\boldsymbol{\omega}_q$	3	∞	∞	3	3	2	3	1	4	14.69
3				$\boldsymbol{\omega}_q$	3	∞	∞	3	3	2	3	1	4	14.69
4			(50,50)	$0.4\boldsymbol{\omega}_q$	2	5	5	∞	∞	4	5	1	7	26.73
5				$0.7\boldsymbol{\omega}_q$	2	5	5	∞	∞	4	5	1	7	32.73
6				$\boldsymbol{\omega}_q$	2	∞	∞	10	10	5	7	2	10	34.25
7		(50,50,50,50)	(5,5)	$0.4\boldsymbol{\omega}_q$	4	2	1	3	5	3	4	1	5	18.23
8				$0.7\boldsymbol{\omega}_q$	4	∞	∞	3	3	3	4	2	5	18.89
9				$\boldsymbol{\omega}_q$	4	∞	∞	3	3	3	4	2	5	18.89
10			(50,50)	$0.4\boldsymbol{\omega}_q$	3	5	5	10	10	5	7	2	9	32.42
11				$0.7\boldsymbol{\omega}_q$	3	5	5	10	10	5	7	2	9	38.42
12				$\boldsymbol{\omega}_q$	4	∞	∞	5	5	4	5	2	7	39.23
13	10	(10,10,10,10)	(5,5)	$0.4\boldsymbol{\omega}_q$	-3	2	2	5	5	3	4	1	5	14.79
14				$0.7\boldsymbol{\omega}_q$	-4	3	3	5	5	3	4	1	6	15.97
15				$\boldsymbol{\omega}_q$	0	∞	∞	3	3	2	3	1	4	16.53
16			(50,50)	$0.4\boldsymbol{\omega}_q$	-3	5	5	∞	20	4	6	1	8	28.16
17				$0.7\boldsymbol{\omega}_q$	-3	5	7	∞	20	4	6	2	10	34.14
18				$\boldsymbol{\omega}_q$	-2	∞	∞	10	10	5	7	2	10	35.64
19		(50,50,50,50)	(5,5)	$0.4\boldsymbol{\omega}_q$	-1	2	2	3	5	3	4	2	6	19.34
20				$0.7\boldsymbol{\omega}_q$	-3	3	3	5	5	4	5	2	7	20.69
21				$\boldsymbol{\omega}_q$	1	∞	∞	3	3	3	4	2	6	20.98
22			(50,50)	$0.4\boldsymbol{\omega}_q$	-1	5	5	10	10	5	7	2	9	34.02
23				$0.7\boldsymbol{\omega}_q$	-1	5	5	10	10	5	7	2	9	40.02
24				$\boldsymbol{\omega}_q$	0	∞	∞	5	10	4	6	3	11	41.23

fixed cost for a time-based dispatch is 40 % of the fixed cost of a quantity-based dispatch, (i.e., $\boldsymbol{\omega}_t = 0.4\boldsymbol{\omega}_q$) the optimal pure time-based policy performs well with an average performance gap of 1.78 % compared to the optimal hybrid policy. In these examples, the optimal pure quantity-based policy performs poorly, with an average cost increase of 15.62 % compared to the hybrid policy. If the relative cost difference between time-based and quantity-based shipments is more moderate, i.e., $\boldsymbol{\omega}_t = 0.7\boldsymbol{\omega}_q$, the quantity-based policy has a better average performance than the time-based policy. However, for individual problems the results are more ambiguous (see Table 3.4). Sometimes, a pure time-based policy performs better than a pure quantity-based policy and vice versa. By definition, the optimal hybrid policy is always best.

3.6 Summary and Outlook

This chapter considers an OWMR inventory distribution system where retailers are clustered into several predefined retailer groups. The retailers face Poisson demand and are replenished by the warehouse. The warehouse receives real-time point-of-sales data from the retailers, why the inventory replenishment process at each retailer is modeled using a continuous review $(S_i - 1, S_i)$ policy. The warehouse uses an (R_0, Q_0) inventory replenishment policy. Shipments from the warehouse to the retailer groups are dispatched according to a hybrid time-and-quantity-based shipment consolidation policy. This means that a shipment is dispatched either when a scheduled time-based shipment day is reached, or a specific consolidation quantity is accumulated.

Our main contribution is the derivation of the PMF of the inventory level at each retailer, which enables the exact evaluation of the expected total inventory, backorder, and shipment costs of the system. Afterwards, we are able to jointly optimize the inventory and shipment policy parameters under this heuristic shipment strategy.

The numerical study demonstrates that there exist scenarios where the dominant hybrid policy can be effectively replaced by either the pure time-based or pure quantity-based consolidation policy. In particular, the pure time-based policy performs very well when the fixed costs for time-based shipments are significantly lower than those for quantity-based shipments, resulting in an average total cost increase of only 1.78 %. Conversely, the optimization of the hybrid policy (T^m, Q^m) leads to the special case of a pure quantity-based policy with $T^m \to \infty$ for all investigated instances when the fixed costs for time-based and quantity-based shipments are equal.

The complexity of our model could be extended by including a compound Poisson demand process or multi-item settings. However, more complex problem settings will increase the computational times for exact solution approaches, making the consideration of heuristics advisable.

3.7 Appendix

3.7.1 Proof of Proposition 1

Proof of Proposition 1. Units corresponding to outstanding orders at retailer i can either be:

(i) Outstanding at the warehouse

(ii) Waiting for a shipment to be dispatched (reserved stock on hand at the warehouse)

(iii) In transit between the warehouse and retailer i

Note that the expected number of units in *(i)* and in *(iii)* are independent of the shipment policy. Consequently, the only change in outstanding orders to retailer i when analyzing a system with the shipment policy compared to a system without the shipment policy is the units corresponding to the reserved stock on hand. As a result

$$E[W_i] = E[O_i|\text{system with the proposed shipment policy}] \\ - E[O_i|\text{system without a shipment policy}]. \tag{3.41}$$

Equation (3.5) then follows from $E[IL_i] = S_i - E[O_i]$. □

3.7.2 Derivation of K^m

In this section, we will derive the PMF of the stochastic variable $K^m \forall\, m \in \{1, 2, \ldots, \mathcal{M}\}$, i.e., the number of units that qualify for shipment to retailer group m between two consecutive time-based dispatches. It is noteworthy that in a system with a hybrid policy, K^m is independent of the consolidation quantity. As a consequence, K^m can be derived by using the exact same method as in a system with a pure time-based consolidation policy.

Stenius et al. (2018) present an exact analysis to obtain the PMF for the number of units that qualify for shipment in between two arbitrary consecutive time-based shipments to retailer group m in a system with a pure time-based consolidation policy. They do not consider the possibility of dispatching quantity-based shipments between the periodic dispatches. However, as all their other modeling assumptions apply to our method, we may directly use their approach to get the distribution of K^m. In this appendix, we recapitulate the main idea of the approach presented in Stenius et al. (2018). For a more detailed explanation, we refer to their paper.

The key to their analysis is to first compute the PMF of the total number of units that have been qualified between two arbitrary consecutive shipments dispatched at time t_0 and $t_1 = t_0 - T^m$, respectively. We denote this quantity by K; thus, the stochastic variable K is the number of units that qualify for shipment to all retailers in the interval $(t_1 = t_0 - T^m, t_0]$. The PMF of the number of qualified units corresponding to retailer group m, K^m, is then obtained by binomial disaggregation.

$$Pr(K^m = k^m) = \sum_{k=k^m}^{\infty} Pr(K = k)\binom{k}{k^m}\left(\frac{\lambda^m}{\Lambda}\right)^{k^m}\left(1 - \frac{\lambda^m}{\Lambda}\right)^{k-k^m} \tag{3.42}$$

It remains to derive $Pr(K = k)$. Similar to Stenius et al. (2018), let us define:

$D_0(s,t)$ Demand at the central warehouse in time interval $(s,t]$, $s < t$, Poisson distributed with mean $\Lambda(t-s)$

$IP_0^-(t)$ Inventory position at the central warehouse at time t, uniformly distributed on $[R_0 + 1, R_0 + Q_0]$

$\underset{R_0,Q_0}{\mathrm{mod}}(x)$ $x + nQ_0$, where n is an integer such that $R_0 < x + nQ_0 \leq R_0 + Q_0$

To determine the number of qualified units to all retailer groups between two consecutive shipments, $t_1 = t_0 - T^m$ and t_0, we also need to take into account the last time instance where the central warehouse can place a replenishment order which will arrive before the respective time-based dispatch. In order for a unit to arrive at the warehouse before t_1 it must have been ordered by the warehouse before $t_1 - L_0$. Similarly, in order for a unit to arrive at the warehouse before t_0 it must have been ordered by the warehouse before $t_0 - L_0$.

These four time instances $t_0, t_1, t_0 - L_0$, and $t_1 - L_0$ form three time intervals. Depending on the length of T^m and L_0, these time instances occur in different sequences. As a consequence, we have to consider two cases: $T^m \leq L_0$ and $T^m > L_0$.

When $T^m \leq L_0$, the final expression to obtain the total number of qualified units between t_1 and t_0 is given by

$$\begin{aligned} K = {} & D_0(t_1, t_0 - L_0) + D_0(t_0 - L_0, t_0) \\ & + \Big(IP_0(t_1 - L_0) - D_0(t_1 - L_0, t_1)\Big)^- \\ & - \Big(\operatorname*{mod}_{R_0, Q_0} \Big(IP_0(t_1 - L_0) - D_0(t_1 - L_0, t_1) \\ & - D_0(t_1, t_0 - L_0)\Big) - D_0(t_0 - L_0, t_0)\Big)^-, \end{aligned} \tag{3.43}$$

whereas if $T^m > L_0$, the expression changes and K is obtained by

$$\begin{aligned} K = {} & D_0(t_1, t_0) + \Big(IP_0(t_1 - L_0) \\ & - D_0(t_1 - L_0, t_0 - L_0) - D_0(t_0 - L_0, t_1)\Big)^- \\ & - \Big(\operatorname*{mod}_{R_0, Q_0} \Big(IP_0(t_1 - L_0) - D_0(t_1 - L_0, t_0 - L_0)\Big) \\ & - D_0(t_0 - L_0, t_1) - D_0(t_1, t_0)\Big)^-. \end{aligned} \tag{3.44}$$

It is valid for both cases that all the stochastic variables are independent of each other. Therefore, the PMF for the total number of units qualified between $(t_1, t_0]$ (i.e., to all retailer groups) can be obtained by convolutions.

3.7.3 Proof of Proposition 2

Proof of Proposition 2. If the considered n^{th} unit is shipped in a potential quantity-based shipment, it is the u^{th} ($u \in \{1, 2, \dots, Q^{\mathcal{M}}\}$) unit to be qualified for shipment among the units in the same shipment. In other words, the last unit in the potential quantity-based shipment containing the considered n^{th} unit is the $(Q^{\mathcal{M}} - u)^{th}$ unit to be qualified for shipment after the *considered* n^{th} *unit*. This means that the priority of the last unit is given by

$$q = n - (Q^{\mathcal{M}} - u). \tag{3.45}$$

It remains to determine u.

We first define $x = D^{\mathcal{M}}(\tau' - T_I, \tau')$, i.e., the demand from retailer group $\mathcal{M}$ between the time-based shipment dispatched at $\tau' - T_I$ and τ'. Moreover, when the time-based shipment is dispatched at $\tau' - T_I$, we assume that there are in total $b_0^{\mathcal{M}}$ units that have been ordered by retailers in retailer group $\mathcal{M}$ before $\tau' - T_I$ that are not included in this shipment due to a stock-out. In other words, these units are backordered at time $\tau' - T_I$.

There are $b_0^{\mathcal{M}} + x$ units that have been ordered before τ' and that qualify for shipment after $\tau' - T_I$. We know that $n \leq b_0^{\mathcal{M}} + x$ since otherwise the considered n^{th} unit would have been dispatched on (or before) the time-based shipment at $\tau' - T_I$. This means that there are $(b_0^{\mathcal{M}} + x - n)$ units with higher priority than the considered n^{th} unit among these $b_0^{\mathcal{M}} + x$ units.

From Figure 3.4 and Figure 3.5 we can see that

$$x = \begin{cases} \alpha_1^A + \alpha_2^A & \textit{for Case A} \\ \alpha_1^B & \textit{for Case B.} \end{cases} \tag{3.46}$$

If we now consider the Euclidian division $(b_0^{\mathcal{M}} + x - n)/Q^{\mathcal{M}}$, the quotient will represent the number of full (potential) quantity-based shipment that are dispatched before the potential quantity-based shipment with the considered n^{th} unit is dispatched. As a consequence, the remainder of the division is the number of units among $(b_0^{\mathcal{M}} + x - n)$ units (with higher priority than the *considered* n^{th} *unit*) that will be in the same potential quantity-based shipment as the *considered* n^{th} *unit*. Since the considered n^{th} unit is the next unit to be qualified for shipment after these units have qualified, we get

$$u = 1 + \mathrm{mod}_{Q^{\mathcal{M}}}(b_0^{\mathcal{M}} + x - n). \tag{3.47}$$

Putting Equation (3.45), Equation (3.46), and Equation (3.47) together we get Equation (3.19) in Proposition 2. □

3.7.4 Derivation of Proposition 4

In this section, we derive the distributions of $X_{\beta_j^i}$ $\forall\, i, j$ presented in Proposition 4. We first need to establish some useful results in Proposition 9 and Corollary 1. Thereafter, we provide the actual derivation of Proposition 4.

Proposition 9. *For a time interval* $(t_0, t_1]$ *assume that we know the number of retailer demands from retailer group* $\mathcal{M}$ *and the other retailer groups respectively, i.e.,* $D_{\mathcal{M}}(t_0, t_1) = \alpha$ *and* $D_{\overline{\mathcal{M}}}(t_0, t_1) = \beta$*. Then the probability that exactly* X *demands from the other retailer groups occur before the* k^{th} *demand from retailer group* $\mathcal{M}$ *is given by a beta-binomial distribution with PMF*

$$Pr(X = x|k, \alpha, \beta) = \binom{\beta}{x} \frac{B(x + k, \beta - x + \alpha - k + 1)}{B(k, \alpha - k + 1)}, \tag{3.48}$$

where $B(m_1, m_2)$ *is the Beta-function defined for positive integers* m_1 *and* m_2 *as*

$$B(m_1, m_2) = \frac{(m_1 - 1)!(m_2 - 1)!}{(m_1 + m_2 - 1)!}. \tag{3.49}$$

We denote the distribution of X *by* $X \sim$ *Betabin*(k, α, β)*.*

Proof of Proposition 9. Properties of the Poisson process assure that the arrival times for a fixed number of demand arrivals during a fixed time interval have a uniform distribution on the considered time interval $(t_0, t_1]$.

As a consequence, the k^{th} demand from the retailers in retailer group $\mathcal{M}$ can be seen as the k-*order statistic* from α uniform random variables. Using properties of the k-order statistic among uniform random variables, we know that the proportion of time until the k^{th} demand occurs, denoted by b, is beta-distributed with parameters

$$b \sim Beta(k, \alpha - k + 1). \tag{3.50}$$

The arrival times for the β demands from the other retailer groups $\overline{\mathcal{M}}$ also have a uniform distribution over the considered time interval $(t_0, t_1]$. Therefore, for a given value of b, the number of demands from the other retailers $\overline{\mathcal{M}}$ that occur before b, is given by a binomial distribution with

$$X \sim Bin(\beta, b). \tag{3.51}$$

Because the parameter b in this distribution is, in fact, a beta-distributed random variable, the distribution of X is given by a binomial distribution with a beta-distributed random parameter. Such distribution is often referred to as a beta-binomial distribution

$$X \sim Betabin(k, \alpha, \beta) \tag{3.52}$$

with PMF

$$Pr(X = x|k, \alpha, \beta) = \binom{\beta}{x} \frac{B(x + k, \beta - x + \alpha - k + 1)}{B(k, \alpha - k + 1)}, \tag{3.53}$$

where $B(m_1, m_2)$ is the Beta-function defined for positive integers m_1 and m_2 as

$$B(m_1, m_2) = \frac{(m_1 - 1)!(m_2 - 1)!}{(m_1 + m_2 - 1)!}. \tag{3.54}$$

□

We also need the following Corollary to Proposition 9.

Corollary 1. *Considering the time interval* (t_{-z}, t_0) *where* t_0 *is an arbitrary point in time, and* t_{-z} *is the time instance when the* z^{th} *system demand before* t_0 *took place. Out of the* $z - 1$ *units demanded in this interval, it is known that* α *demands originate from retailer group* $\mathcal{M}$ *and* β *originate from the other retailer groups* $\overline{\mathcal{M}}$.

Then the probability that exactly $X \leq \beta$ *demands from the other retailer group occur* ***before*** *the* k^{th} ***earliest*** *demand counted backwards from* t_0 *from retailer group* $\mathcal{M}$ *is given by a beta-binomial distribution*

$$X \sim \textit{Betabin}(k, \alpha, \beta), \tag{3.55}$$

where $0 \leq k \leq \alpha$.

Moreover, the probability that exactly $X \leq \beta$ *demands from the other retailer groups occur* ***after*** *the* k^{th} ***most recent*** *demand from retailer group* $\mathcal{M}$ *is given by a beta-binomial distribution*

$$X \sim \textit{Betabin}(k, \alpha, \beta), \tag{3.56}$$

where $0 \leq k \leq \alpha$.

Proof of Corollary 1. When considering a fixed time interval, it is well known that the arrival times of a known number of Poisson arrivals within this time interval, have a uniform distribution on the considered time interval. In Corollary 1, one end-point of the time interval is fixed and the other is determined by the z^{th} demand before the fixed endpoint. By showing that the arrival times of the $z-1$ demands are uniformly distributed between the time of the z^{th} demand and the fixed endpoint, the results follow directly from the proof of Proposition 9.

Let $U_1, U_2, \ldots, U_J$ be J uniform random variables on a fixed arbitrary interval $[t_0, t_{(J+1)}]$. Moreover, we introduce the notations $t_0 \equiv U_{(0)}$ and $t_{(J+1)} \equiv U_{(J+1)}$ for the endpoints of the interval. We also define the order statistics of the J uniform random variables and the endpoints as

$$t_0 \equiv U_{(0)}, U_{(1)}, U_{(2)}, \ldots, U_{(J)}, U_{(J+1)} \equiv t_{(J+1)}. \tag{3.57}$$

It is easy to show that, given the maximum value, $U_{(J)} = t_J$, the first $J-1$ order statistics are distributed as $J-1$ order statistics from $J-1$ uniform random variables in the interval $[t_0, t_J]$. By repeating this argument, we know that given $U_{(k)} = t_k$ the $k-1$ order statistics smaller than $U_{(k)}$ are distributed as the order statistics of $k-1$ uniform random variables on the interval $[t_0, t_k]$. This is true for any $k = 2, 3, \ldots, J$. Due to the symmetry, we also know that $U_{(J-z)} = t_{(J-z)}$ the $z-1$ order statistics larger than $U_{(J-z)}$ are distributed as the order statistics of z uniform random variables on the interval $[t_{J-z}, t_{(J+1)}]$. Moreover, this means that the arrival times of the $z-1$ demands are uniformly distributed between the time of the z^{th} demand and the fixed endpoint t_0 and, thus, the proposition follows directly from Proposition 9. □

Derivation of Proposition 4

In this section, we will focus on determining the distribution of $X_{\beta_j^i}$ from Proposition 4 in Section 3.3.2. We start by recapitulating the definition of $X_{\beta_j^i}$.

$X_{\beta_j^i}$ is the (random) number of demands, among the β_j^i demands from the other retailer groups $\overline{\mathcal{M}}$, that occur in Interval $j \in \{1, 2, 3\}$, before the last unit is ordered by retailer group $\mathcal{M}$. Note that these random variables are only relevant for $q \leq \alpha_1^i + \alpha_2^i + \alpha_3^i$.

$X_{\beta_4^i}$ is the number of demands from the other retailer groups $\overline{\mathcal{M}}$ that have been ordered after the last unit but before $\tau' - T_I - L_0$. Note that these random variables are only relevant for $q > \alpha_1^i + \alpha_2^i + \alpha_3^i$.

We will start by considering $X_{\beta_j^i}$ for $j \in \{1, 2, 3\}$, and then we continue with $X_{\beta_4^i}$. Finally, we summarize the derivation by providing Table 3.2 with distributions and parameters for different cases and situations.

Derivation of $X_{\beta_j^i}$ for $j \in \{1, 2, 3\}$

Let us first consider $X_{\beta_j^i}$ for $j \in \{1, 2, 3\}$. By design, the different values of j represent the time interval in which the retailer demand of the last unit takes place. If the retailer demand takes place in time Interval j, $X_{\beta_j^i}$ denotes how many of the β_j^i demands that take place before the retailer demand of the last unit.

A key insight to determining $X_{\beta_j^i}$ is that b_0 and $b_0^{\mathcal{M}}$ (i.e., the number of backorders at time $\tau' - T_I$) may provide us with information on the order of the demands in an interval from retailer group $\mathcal{M}$ and the other groups $\overline{\mathcal{M}}$, respectively. As we have conditioned the analysis on fixed realizations of the number of backorders, we have to take this into account in the derivation.

The FCFS allocation assures that the demands that are backordered at $\tau' - T_I$ are the demands that have occurred most recently before this time. As a consequence, we know that out of the b_0 most recent demands to have occurred before $\tau' - T_I$, $b_0^{\mathcal{M}}$ must have been demanded by retailer group $\mathcal{M}$ (and $b_0 - b_0^{\mathcal{M}}$ have been demanded by the other retailer groups $\overline{\mathcal{M}}$). In situations where some of the demands during the interval are backordered, and some are not backordered, we need to take into consideration that b_0 and $b_0^{\mathcal{M}}$ will influence how many of the β_j^i demands from the retailers in $\overline{\mathcal{M}}$ that occurred before the demand for the last unit, i.e., the random variable $X_{\beta_j^i}$.

To proceed with the analysis, we will distinguish between three different situations:

Situation 1 The demands that have occurred in the same interval as the last unit are **all** backordered at time $\tau' - T_I$

Situation 2 **None** of the demands that have occurred in the same interval as the last unit are backordered at time $\tau' - T_I$.

Situation 3 Some of the demands that have occurred in the same interval as the last unit are backordered at time $\tau' - T_I$ and some are not.

Situation 1 and Situaion 2 are the most straightforward to consider, i.e., when either **all** demands in the same interval as the last unit was demanded are backordered at $\tau' - T_I$, or **none** of the demands in the same interval are backordered at $\tau' - T_I$. It is worth noting that the latter is always true if the last unit was ordered after $\tau' - T_I$. In these two situations, knowing that $b_0^{\mathcal{M}}$ out of the b_0 backorders belong to retailer group $\mathcal{M}$ does not provide any additional information. We may, therefore, directly apply Proposition 9 to obtain the distribution of $X_{\beta_j^i}$. This is done by letting the last unit be the k^{th} demand from retailer group $\mathcal{M}$ counted from the beginning of the time interval that the demand occurs and replacing α and β in Equation (3.48) with α_j^i and β_j^i for the corresponding interval and case. To conclude, in the special case where the system demands in Interval j (i.e., the interval in which the last unit is demanded) are either **all** backordered at time $\tau' - T_I$ or **none** are backordered at this time, the distribution of $X_{\beta_j^i}$ is given by

$$X_{\beta_j^i} \sim Betabin(k, \alpha_j^i, \beta_j^i), \tag{3.58}$$

where k depend on which interval the last unit was demanded in, i.e.,

$$k = \begin{cases} \alpha_1^i - q + 1 & 0 < q \leq \alpha_1^i \\ \alpha_1^i + \alpha_2^i - q + 1 & \alpha_1^i < q \leq \alpha_1^i + \alpha_2^i \\ \alpha_1^i + \alpha_2^i + \alpha_3^i - q + 1 & \alpha_1^i + \alpha_2^i < q \leq \alpha_1^i + \alpha_2^i + \alpha_3^i. \end{cases} \tag{3.59}$$

To exemplify, in Case A, if $k = 4$ and the last unit is demanded in Interval 2, i.e. in $(\tau' - T_I, \tau' - L_0]$. The last unit is the fourth unit to be demanded counting from $\tau' - T_I$. It should be noted that the three different cases in the expression for k Equation (3.59) simply determine the time interval in which retailer group $\mathcal{M}$ demands the last unit, see Figure 3.4 and Figure 3.5 in Section 3.3.1. Moreover, $X_{\beta_2^A}$ is then the number of demands among the β_2^A demands that occur before the fourth demand from retailer group $\mathcal{M}$ occurs.

We will now focus on Situation 3, i.e., when some of the system demands in the same time interval as the last unit is demanded are backordered at time $\tau' - T_I$ and some are **not** backordered, i.e., are dispatched with (or before) the time-based shipment. For Case A, this situation may only occur if the last unit is demanded in Interval 3. Similarly, for the situation to occur in Case B, the last unit must have been demanded in either Interval 2 or Interval 3. Note that if the last unit is demanded by retailer group $\mathcal{M}$ after $\tau' - T_I$, none of the retailer demands in the same time interval can be backordered at time $\tau' - T_I$, thus Situation 2 from above applies. See Figure 3.4 and Figure 3.5.

Moreover, as we only need to consider the situation where the considered n^{th} unit is backordered (otherwise it would have been dispatched on a time-based shipment), and the last unit is demanded after the considered n^{th} unit, we know that the last unit is also backordered (or not yet demanded at time $\tau' - T_I$). As a consequence, we know that all the demands from the other retailers that are **not** backordered at time $\tau' - T_I$ must have been demanded before the last unit. We denote this demanded quantity by ξ. We start by determining ξ.

$$\xi = \begin{cases} \beta_3^A - (b_0 - b_0^{\mathcal{M}}) & \text{for Case A } \textbf{and } \alpha_1^A + \alpha_2^A < q \leq \alpha_1^A + \alpha_2^A + \alpha_3^A \\ \beta_2^B - (b_0 - b_0^{\mathcal{M}}) & \text{for Case B } \textbf{and } \alpha_1^B < q \leq \alpha_1^B + \alpha_2^B \\ \beta_2^B + \beta_3^B - (b_0 - b_0^{\mathcal{M}}) & \text{for Case B } \textbf{and } \alpha_1^B + \alpha_2^B < q \leq \alpha_1^B + \alpha_2^B + \alpha_3^B \end{cases} \tag{3.60}$$

Note that the value of ξ is a known, deterministic value for given values of b_0, $b_0^{\mathcal{M}}$, $\boldsymbol{\alpha^i}$, and $\boldsymbol{\beta^i}$. The remaining part of $X_{\beta_j^i}$ (i.e., of the units that are ordered before the last unit in the same interval) are units among the backordered demands from the other retailer groups $\overline{\mathcal{M}}$ that have been ordered before the last unit. We use Corollary 1 to get this quantity.

To apply Corollary 1, we identify the most recent demand out of all the demands that are **not** backordered at time $\tau' - T_I$. Moreover, this demand takes place in the same time interval as the last unit (otherwise we would not be in Situation 3). Following the notations from Corollary 1, this time interval ends with the time instance t_0. We know that

$$t_0 = \begin{cases} \tau' - T_I & \text{for Case A } \textbf{and } \alpha_1^A + \alpha_2^A < q \leq \alpha_1^A + \alpha_2^A + \alpha_3^A \\ \tau' - T_I & \text{for Case B } \textbf{and } \alpha_1^B < q \leq \alpha_1^B + \alpha_2^B \\ \tau' - L_0 & \text{for Case B } \textbf{and } \alpha_1^B + \alpha_2^B < q \leq \alpha_1^B + \alpha_2^B + \alpha_3^B. \end{cases} \tag{3.61}$$

Using the notation, z, from Corollary 1, the most recent demand out of all the demands that are **not** backordered at time $\tau' - T_I$, is the z^{th} most recent demand to have occurred prior to t_0 (note that z counts the total demand in the system and not only the demand at retailer group $\mathcal{M}$).

Moreover, α and β are the backordered demands (from retailer group $\mathcal{M}$ and the other retailer groups respectively) that were demanded in the same interval as the last unit (note that $z = \alpha + \beta + 1$). Thus, using these notations and Corollary 1, we have

$$X_{\beta_j^i} = \xi + X \tag{3.62}$$

where

$$X \sim Betabin(k, \alpha, \beta) \tag{3.63}$$

and ξ is given by Equation (3.60). α and β take different values depending on which interval the last unit is demanded in and which of Case A and Case B that we consider. Recall that for Case A, the situations that we consider can only happen if $\alpha_1^A + \alpha_2^A < q \leq \alpha_1^A + \alpha_2^A + \alpha_3^A$. Thus, the α and β parameters for the beta-binomial distribution equals all the backordered demands at time $\tau' - T_I$ for retailer group $\mathcal{M}$ and the other retailer groups respectively, i.e.,

$$\alpha = b_0^{\mathcal{M}} \qquad \alpha_1^A + \alpha_2^A < q \leq \alpha_1^A + \alpha_2^A + \alpha_3^A \tag{3.64}$$

$$\beta = b_0 - b_0^{\mathcal{M}} \qquad \alpha_1^A + \alpha_2^A < q \leq \alpha_1^A + \alpha_2^A + \alpha_3^A. \tag{3.65}$$

For Case B, the last unit can be demanded in both the second and the third interval; thus,

$$\alpha = \begin{cases} b_0^{\mathcal{M}} & \alpha_1^B < q \leq \alpha_1^B + \alpha_2^B \\ b_0^{\mathcal{M}} - \alpha_2^B & \alpha_1^B + \alpha_2^B < q \leq \alpha_1^B + \alpha_2^B + \alpha_3^B \end{cases} \tag{3.66}$$

$$\beta = \begin{cases} b_0 - b_0^{\mathcal{M}} & \alpha_1^B < q \leq \alpha_1^B + \alpha_2^B \\ b_0 - b_0^{\mathcal{M}} - \beta_2^B & \alpha_1^B + \alpha_2^B < q \leq \alpha_1^B + \alpha_2^B + \alpha_3^B. \end{cases} \tag{3.67}$$

Moreover, the k^{th} unit is only counted among the backordered demand from retailer group $\mathcal{M}$ in this interval. The value of k also depends on which interval the last unit is ordered in

$$k = \begin{cases} \alpha_1^i + b_0^{\mathcal{M}} - q + 1 & \alpha_1^i < q \le \alpha_1^i + \alpha_2^i \\ \alpha_1^i + \alpha_2^i + b_0^{\mathcal{M}} - q + 1 & \alpha_1^i + \alpha_2^i < q \le \alpha_1^i + \alpha_2^i + \alpha_3^i. \end{cases} \tag{3.68}$$

Derivation of $X_{\beta_4^i}$

We now focus on $X_{\beta_4^i}$, i.e., when the last unit is demanded by retailer group $\mathcal{M}$ prior to $\tau' - T_I - L_0$ (before Interval 3). Since the considered n^{th} unit is backordered at time $\tau' - T_I$, the last unit must also be backordered at this time. In this case, $X_{\beta_j^i}$ determines how many of the demands from the other retailer groups $\overline{\mathcal{M}}$ that are demanded **after** the last unit, but before $\tau' - T_I - L_0$. Note that there is not a closed interval in this case. However, in line with the previous notations, we refer to this as $X_{\beta_4^i}$, i.e., $j = 4$, since it occurs before Interval 3.

Corollary 1 can be directly applied to this situation. The notations correspond to $t_0 = \tau' - T_I - L_0$ and $z - 1$ as the total number of demands that have occurred prior to t_0 and are still backordered at time $\tau' - T_I$. The parameters α and β are the number of demands, from retailer group $\mathcal{M}$ and the other retailers $\overline{\mathcal{M}}$ respectively, that have occurred before $\tau' - T_I - L_0$ and are still backordered at time $\tau' - T_I$. This means that

$$\alpha = \begin{cases} b_0^{\mathcal{M}} - \alpha_3^A & \text{for Case A} \\ b_0^{\mathcal{M}} - \alpha_3^B - \alpha_2^B & \text{for Case B} \end{cases} \tag{3.69}$$

$$\beta = \begin{cases} (b_0 - b_0^{\mathcal{M}}) - \beta_3^A & \text{for Case A} \\ (b_0 - b_0^{\mathcal{M}}) - \beta_3^B - \beta_2^B & \text{for Case B.} \end{cases} \tag{3.70}$$

Moreover, the last unit is the k^{th} most recent demand at retailer group $\mathcal{M}$ to occur before $\tau' - T_I - L_0$, i.e., in both Case A and Case B, we have

$$k = q - \alpha_3^i - \alpha_2^i - \alpha_1^i. \tag{3.71}$$

We summarize the above derivation of $X_{\beta_j^i}$ in Table 3.2 by letting $X_{\beta_j^i} = \xi + X$, where $X \sim Betabin(k, \alpha, \beta)$ with the parameters given in Table 3.2.

3.7.5 Proof of Proposition 5

Proof of Proposition 5. Recall that the analysis is conditioned on $IP_0(\tau' - T_I - L_0) = ip$. Moreover, the retailer demand that occurs at time t is the ψ^{th} demand after $\tau' - T_I - L_0$, i.e., $D_0(\tau' - T_I - L_0, t) = \psi$. Let $IP_0(t)$ denote the inventory position at the central warehouse just before t.

Let us first consider y such that $ip = R_0 + Q_0 - y$, i.e.,

$$y = R_0 + Q_0 - ip, \tag{3.72}$$

where y represents the number of demands that have occurred since the inventory position most recently was equal to $R_0 + Q_0$. As $IP_0(t) \in \{R_0 + 1, R_0 + 2, \ldots, R_0 + Q_0\}$, we have $y \in \{0, 1, \ldots, Q_0 - 1\}$.

$IP_0(t)$ represents the warehouse inventory position after an additional $\psi - 1$ demands from the retailers. As a consequence,

$$\begin{aligned} IP_0(t) &= R_0 + Q_0 - \mathrm{mod}_{Q_0}(y + (\psi - 1)) \\ &= R_0 + Q_0 - \mathrm{mod}_{Q_0}(R_0 + Q_0 - ip + (\psi - 1)) \\ &= R_0 + Q_0 - \mathrm{mod}_{Q_0}(R_0 - ip + (\psi - 1)). \end{aligned} \tag{3.73}$$

Just before the retailer order at time t takes place the inventory position is $IP_0(t)$. By definition of ϕ, the warehouse order including the unit that will satisfy this retailer order takes place $\phi(\psi|ip)$ system demands prior to t. As $IP_0(t)$ is known from Equation (3.73) we want to determine ϕ for a given $IP_0(t)$.

Let x be the number of demands prior to t when the most recent warehouse order took place, i.e.,

$$IP_0(t) = R_0 + Q_0 - x. \tag{3.74}$$

This means that the second most recent warehouse order took place $x + Q_0$ demands prior to t and the third most recent warehouse order took place $x + 2Q_0$ demands prior to t. In general, the k^{th} most recent warehouse order took place $x + (k-1)Q_0$ demands prior to t. Note that $k = 0$ if the warehouse order took place after t (this situation is only relevant for $R_0 < 0$).

After the warehouse order takes place, the unit will satisfy the demands occurring $R_0+1, R_0+2, \dots, R_0+Q_0$ demands later. As a consequence, the unit that satisfies a demand at time t must have been ordered $\phi = R_0+1+y$ demands earlier where $y = 0, 1, \dots, Q_0-1$. Since this warehouse order took place $x+(k-1)Q_0$ demands earlier (or later in case $k=0$), we know that $R_0+1+y = x+(k-1)Q_0$. This means that

$$y = x - R_0 - 1 + (k-1)Q_0 \tag{3.75}$$

where k is chosen such that $y = 0, 1, \dots Q_0-1$. Since $x = R_0+Q_0-ip$ we have

$$y = kQ_0 - ip, \tag{3.76}$$

where again n is chosen such that $y = 0, 1, \dots Q_0-1$.

Thus,

$$\begin{aligned} \phi(\psi|ip) &= R_0 + 1 + y \\ &= R_0 + 1 + \mathrm{mod}_{Q_0}(-ip) \\ &= R_0 + Q_0 - \mathrm{mod}_{Q_0}(ip-1). \end{aligned} \tag{3.77}$$

We plug in Equation (3.73) into Equation (3.77) and simplify the expression

$$\begin{aligned} \phi(\psi|ip) &= R_0 + Q_0 - \mathrm{mod}_{Q_0}\left(R_0 + Q_0 - \mathrm{mod}_{Q_0}(R_0 - ip + (\psi-1)) - 1\right) \\ &= R_0 + Q_0 - \mathrm{mod}_{Q_0}\left(R_0 - 1 - \mathrm{mod}_{Q_0}(R_0 - 1 - ip + \psi)\right) \\ &= R_0 + Q_0 - \mathrm{mod}_{Q_0}\left(ip - \psi\right) \end{aligned} \tag{3.78}$$

□

3.7.6 Proof of Proposition 6

Proof of Proposition 6. In Case A, Interval 2, i.e., $(\tau' - T_I, \tau' - L_0]$ depends on the random variable T_I. As a consequence, the length of Interval 2 is uniformly distributed on $(0, T^{\mathcal{M}} - L_0)$, whereas all the other time intervals have a fixed length of L_0 (see Figure 3.4 for details). Let X be the random length of Interval 2; this means that

$$D^A(\tau' - T_I, \tau' - L_0) \sim Poisson(\Lambda X), \text{ where } X \sim Uniform(0, T^{\mathcal{M}} - L_0). \tag{3.79}$$

Thus, the number of demands in this interval is distributed according to a Poisson distribution with a uniformly distributed parameter. The objective is to derive the PMF of this distribution, i.e., $\mathcal{P}_A(\gamma_2^A) = Pr(D^A(\tau' - T_I, \tau' - L_0) = \gamma_2^A)$. By using the law of total probability, we get

$$\begin{aligned}
\mathcal{P}_A(\gamma_2^A) &= \int_0^{T^{\mathcal{M}}-L_0} Pr(D^A(\tau' - T_I, \tau' - L_0) = \gamma_2^A) Pr(T_I - L_0 = x) dx \\
&= \frac{1}{T^{\mathcal{M}} - L_0} \int_0^{T^{\mathcal{M}}-L_0} \frac{(\Lambda x)^{\gamma_2^A}}{\gamma_2^A!} \mathrm{e}^{-\Lambda x} dx \\
&= \frac{1}{T^{\mathcal{M}} - L_0} \frac{\Lambda^{\gamma_2^A}}{\gamma_2^A!} \int_0^{T^{\mathcal{M}}-L_0} x^{\gamma_2^A} \mathrm{e}^{-\Lambda x} dx. \qquad (3.80)
\end{aligned}$$

The latter integral can be solved analytically by repeated use of integration by parts (Gradštejn et al., 1994). The result has the form

$$\int_0^b x^n \mathrm{e}^{-ax} dx = \frac{n!}{a^{n+1}} \left[1 - \mathrm{e}^{-ab} \sum_{i=0}^{n} \frac{(ab)^i}{i!} \right]. \qquad (3.81)$$

By further simplifying Equation (3.80), using Equation (3.81) and $b = T^{\mathcal{M}} - L_0$, we get

$$\begin{aligned}
\mathcal{P}_A(\gamma_2^A) &= \frac{1}{b} \frac{\Lambda^{\gamma_2^A}}{\gamma_2^A!} \left(\frac{\gamma_2^A!}{\Lambda^{\gamma_2^A+1}} \left[1 - \mathrm{e}^{-\Lambda b} \sum_{i=0}^{\gamma_2^A} \frac{(\Lambda b)^i}{i!} \right] \right) \\
&= \frac{1}{\Lambda b} \left[1 - \mathrm{e}^{-\Lambda b} \sum_{i=0}^{\gamma_2^A} \frac{(\Lambda b)^i}{i!} \right] \\
&= \frac{1}{\Lambda b} \left[1 - F^{\Lambda b}(\gamma_2^A) \right], \qquad (3.82)
\end{aligned}$$

where $F^{\lambda}(k) = \mathrm{e}^{-\lambda} \sum_{j=0}^{k} \frac{\lambda^j}{j!}$. □

3.7.7 Proof of Proposition 7

Proof of Proposition 7. Focusing on Case B, the reader should remember that all three time intervals depend on the random variable T_I (see Figure 3.5). Moreover, since we condition on $T_I \leq L_0$, we know that $T_I \sim Uniform\,(0, b)$ where $b = \min(L_0, T^{\mathcal{M}})$. To derive the proposition, we condition on $T_I = x$ and take the expectation across all possible values of x, i.e., $[0, b]$. This means that Interval 1 and Interval 3 are of length x and Interval 2 is of length $L_0 - x$, see Figure 3.5 for details. With a fixed interval length, the number of demands in the respective time interval are independent Poisson random variables. Thus we have

$$\begin{aligned}
&\mathcal{P}_B(\gamma_1^B,\gamma_2^B,\gamma_3^B)\\
&= Pr(D^B(\tau'-T_I-L_0,\tau'-L_0)=\gamma_3^B), D^B(\tau'-L_0,\tau'-T_I)=\gamma_2^B), D^B(\tau'-T_I,\tau')=\gamma_1^B)\\
&= \int_0^b \frac{1}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}(\lambda x)^{\gamma_1^B+\gamma_3^B}\mathrm{e}^{-2\lambda x}\left(\lambda(L_0-x)\right)^{\gamma_2^B}\mathrm{e}^{-\lambda(L_0-x))}Pr(T_I=x)dx\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\int_0^b x^{\gamma_1^B+\gamma_3^B}(L_0-x)^{\gamma_2^B}\mathrm{e}^{-2\lambda x-\lambda(L_0-x)}dx\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\int_0^b x^{\gamma_1^B+\gamma_3^B}(L_0-x)^{\gamma_2^B}\mathrm{e}^{-\lambda x}\mathrm{e}^{-\lambda L_0}dx\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\mathrm{e}^{-\lambda L_0}\int_0^b x^{\gamma_1^B+\gamma_3^B}\mathrm{e}^{-\lambda x}\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}L_0^{\gamma_2^B-k}(-1)^k x^k dx\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\mathrm{e}^{-\lambda L_0}\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}L_0^{\gamma_2^B-k}(-1)^k\int_0^b x^{\gamma_1^B+\gamma_2^B+k}\mathrm{e}^{-\lambda x}dx\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\mathrm{e}^{-\lambda L_0}\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}L_0^{\gamma_2^B-k}(-1)^k\frac{(\gamma_1^B+\gamma_3^B+k)!}{\lambda^{\gamma_1^B+\gamma_3^B+k+1}}\left[1-\mathrm{e}^{-\lambda b}\sum_{i=0}^{\gamma_1^B+\gamma_3^B+k}\frac{(\lambda b)^i}{i!}\right]\\
&= \frac{1}{b}\frac{\lambda^{\gamma_1^B+\gamma_2^B+\gamma_3^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\mathrm{e}^{-\lambda L_0}\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}L_0^{\gamma_2^B-k}(-1)^k\frac{(\gamma_1^B+\gamma_3^B+k)!}{\lambda^{\gamma_1^B+\gamma_3^B+k+1}}\left[1-F^{\lambda b}(\gamma_1^B+\gamma_3^B+k)\right]\\
&= \frac{1}{b}\frac{\lambda^{\gamma_2^B}}{\gamma_1^B!\gamma_2^B!\gamma_3^B!}\mathrm{e}^{-\lambda L_0}\sum_{k=0}^{\gamma_2^B}\binom{\gamma_2^B}{k}L_0^{\gamma_2^B-k}(-1)^k\frac{(\gamma_1^B+\gamma_3^B+k)!}{\lambda^{k+1}}\left[1-F^{\lambda b}(\gamma_1^B+\gamma_3^B+k)\right],
\end{aligned} \tag{3.83}$$

where $F^{\lambda}(k)=\sum_{j=0}^{k}\frac{\lambda^j\mathrm{e}^{-\lambda}}{j!}$.

Note that the integral on line 7 is essentially the same as Equation (3.81). □

3.7.8 Proof of Proposition 8

Proof of Proposition 8. First, recall that $B_0(\tau'-T_I)=b_0$ is directly obtained from ip and the demand during $(\tau'-T_I-L_0,\tau'-T_I]$, i.e., the demand during the lead-time from the outside supplier. As the lead-time demand is obtained from $\boldsymbol{\alpha^i}$ and $\boldsymbol{\beta^i}$, we have

$$b_0=\begin{cases}\left(ip-\alpha_3^A-\beta_3^A\right)^- & \text{for Case A}\\ \left(ip-\alpha_2^B-\alpha_3^B-\beta_2^B-\beta_3^B\right)^- & \text{for Case B.}\end{cases} \tag{3.84}$$

Deriving the deterministic part of the expression in Proposition 8, i.e., κ, is straightforward as it corresponds to intervals where all the demands are backordered. For example, let us consider Case B. If $b_0 > \alpha_2^B + \beta_2^B + \alpha_3^B + \beta_3^B$, we know that all demands for both retailer group $\mathcal{M}$ and the other retailer groups $\mathcal{M}$, in both Interval 2 and Interval 3 are backordered. In particular, $b_0^{\mathcal{M}} > \alpha_2^B + \alpha_3^B$, and thus we can focus on determining how many additional backorders there are for retailer group $\mathcal{M}$ except for the $\alpha_2^B + \alpha_3^B$ backorders.

Next, we consider the stochastic part of the expression, i.e., the number of backorders for retailer group $\mathcal{M}$ in intervals where not all the demands are backordered. To proceed with the analysis, we first note that the b_0 backorders have different priorities according to the FCFS policy. The backorder with the highest priority corresponds to the earliest demand that was not satisfied, which is among the demands backordered at time $\tau' - T_I$. Looking backward in time from $\tau' - T_I$, this demand was the b_0^{th} most recent demand to occur prior to $\tau' - T_I$. The remaining derivation hinges upon whether the b_0^{th} demand prior to $\tau' - T_I$ takes place before or after $\tau' - T_I - L_0$. In the context of Case A, this translates to a comparison between b_0 and $\alpha_3^A + \beta_3^A$ to determine if it is smaller or larger. Conversely, for Case B, we compare b_0 to $\alpha_3^A + \alpha_2^A + \beta_3^A + \beta_2^A$.

If the b_0^{th} demand takes place prior to $\tau' - T_I - L_0$ (i.e., before the third interval as defined in Figure 3.4 and Figure 3.5) we need to determine how many of the backordered units demanded before $\tau' - T_I - L_0$ that belong to retailer group $\mathcal{M}$ and to the other retailer groups $\overline{\mathcal{M}}$ respectively. Looking at these backordered units one at a time, the independent Poisson demand processes and the FCFS ensure that the destination for each backordered unit is independent of each other. Thus, a unit is destined to retailer group $\mathcal{M}$ with probability $\lambda^{\mathcal{M}}/\Lambda$ and to the other retailer groups $\overline{\mathcal{M}}$ with probability $1 - \lambda^{\mathcal{M}}/\Lambda$. Consequently, the probability that k out of y backordered units (demanded prior to $\tau' - T_I - L_0$) are destined to retailer group $\mathcal{M}$, is obtained from the binomial distribution $Bin(y, \lambda^{\mathcal{M}}/\Lambda)$.

If the b_0^{th} demand takes place after $\tau' - T_I - L_0$, we know that this demand occurs in one of the previously defined intervals (see Figure 3.4 and Figure 3.5). Thus, it is necessary to determine how many of the backorders, in the same interval as the b_0^{th} demand takes place, that belong to retailer group $\mathcal{M}$ and to the other retailer groups $\overline{\mathcal{M}}$ respectively. If the b_0^{th} demand takes place in Interval j, we know that given Case i, there are α_j^i demands in the interval from retailer group $\mathcal{M}$ and β_j^i demands from the other retailer groups $\overline{\mathcal{M}}$. Following from the properties of independent Poisson demand and FCFS, we know that the sequential order of these α_j^i and β_j^i demands is random and can therefore be regarded as a sample of units with two features, retailer group $\mathcal{M}$ and the other retailer groups $\overline{\mathcal{M}}$, from a finite population without replacement. Thus, the probability that x out of the backordered demands in this interval belongs to retailer group $\mathcal{M}$ is given by the hypergeometric distribution $Hyp(\alpha_j^i + \beta_j^i, \alpha_j^i, b^*)$, where b^* is the number of backordered demands in the j^{th} interval.

□

Chapter

4 Heuristic Time-based Shipment Policy with Flexible Deliveries

The following chapter is based on Ralfs and Kiesmüller (2022). This chapter primarily focuses on inventory control with the integration of shipment consolidation and ADI. We analyze the value of incorporating ADI in a single-echelon inventory system when external orders are satisfied under a flexible time-based shipment consolidation policy. Further, we investigate the influence on the optimal inventory and shipment policy parameters. In Section 4.1, we start with a detailed formulation of the general problem. We consider a continuous-time single-echelon inventory system where outbound shipments are dispatched according to a time-based shipment scheme. ADI is modeled as a demand lead time, and flexible deliveries are allowed, meaning that orders may be fulfilled ahead of their due date. We present the analysis of the shipment and inventory costs of the inventory system in Section 4.2. The computation of the shipment costs is based on convolutions in combination with a procedure to approximate a specific random variable. The exact inventory costs can be obtained based on a unit-tracking approach. Section 4.3 explains how the joint optimization of inventory and shipment parameters is performed. Moving on to Section 4.4, we not only validate the proposed approximate solution approach but also present managerial insights and results. This section highlights the necessity not only to optimize safety stocks separately but also to incorporate and optimize other processes, such as transportation, to achieve higher cost reductions. Lastly, Section 4.5 summarizes the main contributions and findings of this chapter.

4.1 Problem Formulation

We consider a single-item continuous review inventory system composed of one warehouse, which receives random orders from a production facility both belonging to the same company. The warehouse supplies this production facility, which orders according to a Poisson process with rate λ. Note that an extension to several production facilities within one retailer group and similar transportation times is straightforward due to the properties of the Poisson process. The inventory at the warehouse is replenished from an outside supplier with ample capacity and a constant supply lead time L_s according to an (R, Q) policy, indicating that a replenishment order of Q units is placed if the modified inventory position (= stock on hand + outstanding replenishment orders − backorders − observed orders) reaches the reorder level R. The inventory manager has to determine the value of R, and we assume that Q is fixed due to a contract with the outside supplier. As an approximation, Q can also be predetermined by applying the EOQ formula, which has been shown to be close to optimal in previous research (Axsäter, 2015). Therefore, Q is not a decision variable in our model.

The company uses a preorder strategy, and therefore, the production facility is forced to place each order L_d time units before the actual demand occurs, indicating that each order is combined with a due date. The time between order placement and the due date is known as the demand lead time L_d. For clarity, we call orders that have not reached the due date *orders* and orders that have reached the due date *demands*. At the moment when the due date is reached, the warehouse is obligated to satisfy this demand with the next delivery if sufficient stock on hand is available. Companies accept such a preorder strategy when they receive a bonus, e.g., in the form of lower unit costs. Since the warehouse and the facility belong to the same company, we do not include the bonus in our model. This also makes it possible to determine the pure value of inserting ADI. We further assume a constant and identical demand lead time for all production facility orders. Moreover, we focus on perfect order information, meaning that the order quantity and due date are certain and cannot be changed after order placement. Order cancellations may occur, but due to associated high cancellation costs, we assume that this happens very seldom, and we exclude imperfect ADI for the sake of simplicity. In this chapter, we investigate the case in which $L_d \leq L_s$ (make-to-stock) because otherwise there is no need to keep stock at the warehouse (make-to-order).

In contrast to other studies that consider ADI, the warehouse does not satisfy demands immediately but applies a time-based shipment consolidation policy. After a fixed time period T, referred to as the *shipment interval*, a load with all accumulated demands is dispatched from the warehouse to the production facility. In the case of stock-outs, demand is backordered on the shipment day. When sufficient inventory is available again, backorders are satisfied with the next scheduled shipment. Time-based shipping strategies are popular in the industry because scheduling, administration, and coordination of processes at the production facility are easy to manage.

Due to a stochastic order process, it is evident that the number of accumulated demands during a shipment interval, resulting in shipment quantity M, is also random. The warehouse does not have its own fleet of trucks and, therefore, engages a third-party logistics service provider (3PL) for the transportation from the warehouse to the production facility. Although 3PL can react quickly to requests, this flexible strategy is quite costly. To reduce costs, companies negotiate contracts in which a fixed transportation capacity is reserved for a lower price. Thus, we consider a *primary* and an *alternative transportation option*. The primary transportation option reflects the capacity reserved at the 3PL for periodic shipments from the warehouse to the production facility, which is limited to a self-chosen capacity reservation Cap and can be extended by the alternative option, which has ample capacity but is more expensive. The alternative option can only be used on scheduled shipment days and if the reserved capacity is exhausted. For the primary option, fixed costs $\omega(Cap)$ per shipment occur, which depend on the chosen reserved transportation capacity. The fixed shipment costs represent nothing more than variable shipment costs that have to be paid for each reserved transportation capacity unit $(c_1 \cdot Cap)$. However, these costs are charged independently of the realized shipment quantity. If the realized shipment quantity exceeds the reserved capacity of the primary transportation option, the warehouse has to pay variable costs c_2 for each unit shipped by the alternative option. Note that $c_2 > c_1$ to reflect that the alternative transportation is more expensive than a capacity reservation per unit. Dispatching, sorting, and consolidating costs are included in the before-mentioned shipment costs paid to the 3PL. As we focus on the inventory and transportation decisions at the warehouse, and demand has to be satisfied, the transportation time from the warehouse to the production facility is not relevant to the decision and, therefore, is not included in the model.

Considering that the production facility orders in advance, the question arises regarding how to include information about future demand in the transportation schedule. It is allowed that orders can be satisfied before their due date is reached, known as flexible delivery. In our approach, demand must be shipped to the production facility on the upcoming shipment day, if necessary, using the alternative transportation option. In contrast, orders can only be dispatched if there is remaining capacity at the primary transportation option. However, it is only allowed to ship an order one shipment date ahead because the production facility cannot keep a large amount of additional stock.

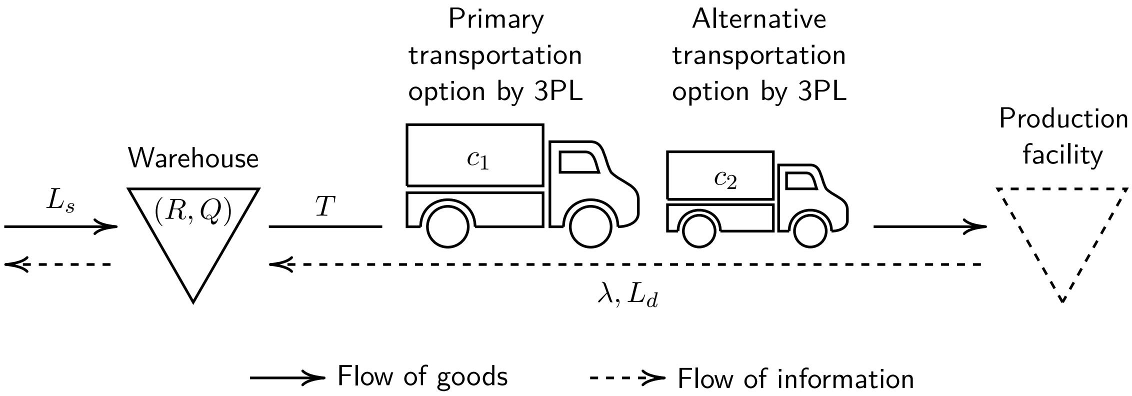

Figure 4.1: Flow of information and goods in the considered inventory system

The application of this approach results in increased utilization of the reserved capacity. At the same time, the usage of the expensive transportation option can be reduced if the reserved capacity is not fully exhausted on average. At the warehouse, orders and demands are allocated according to the FCFS principle.

In addition to the shipment costs, three other types of costs occur and influence the decision-making. First, inventory holding costs h are charged for each unit on stock per time unit. Second, late-delivery costs l arise for each unit that waits due to the shipment consolidation policy per time unit. These costs are related to a discount for the production facility since it must hold more safety stock to reach the same service level. Backorder costs at the warehouse are included in the late-delivery costs because backorders lead to a longer waiting time. Third, costs for early deliveries are considered, compensating the facility that requires more space to store units shipped prior to the due date. Early-delivery costs e are charged for each early shipped unit per time unit. In the following, we use the term *inventory costs* for the sum of these three types of costs, which are assumed to be linear. An overview of the problem is presented in Figure 4.1.

We assume a central control of the system where the inventory manager has to determine the shipment interval as well as the amount of safety stock needed at the warehouse. He or she can further determine in advance how much transportation capacity to reserve for a low price. To summarize, the system's expected total cost per time unit $TC(R, T, Cap)$ is composed of the expected inventory costs $TIC(R, T, Cap)$ and expected shipment costs $TSC(R, T, Cap)$.

$$TC(R, T, Cap) = TIC(R, T, Cap) + TSC(R, T, Cap) \tag{4.1}$$

A low stock level at the warehouse leads to high late-delivery and shipment costs because orders can hardly ever be shipped in advance; thus, the flexible delivery option cannot be used at all. However, a high stock level would cause high inventory holding and early-delivery costs at the warehouse since there is always sufficient stock to dispatch orders before the due date. Conversely, a small shipment interval results in low late-delivery and early-delivery costs but leads to high shipment costs. The right balance between these types of costs must be found; therefore, the following optimization problem is formulated:

$$\min TC(R,T,Cap) \qquad R \in \mathbb{Z}, T \in \mathbb{N}, Cap \in \mathbb{N}_0 \tag{4.2}$$

To determine the optimal policy parameters R^*, T^*, and Cap^*, it is necessary to be able to evaluate a policy; therefore, mathematical expressions for the expected total cost must be derived, which are the focus of the next section.

4.2 Analysis

In this section, we derive expressions for the expected shipment costs $TSC(R,T,Cap)$ per time unit and the expected inventory costs $TIC(R,T,Cap)$ per time unit. Our analysis of the latter relies on the unit tracking methodology introduced by Axsäter (1990), where each unit going through the system is observed separately, which enables us to calculate the related expected costs for each unit. However, we must adapt this methodology for a situation with ADI, limited transportation capacity, and flexible deliveries.

4.2.1 Expected Shipment Costs

First, we focus on the expected shipment costs per time unit, where fixed costs $\omega(Cap)$ occur for each shipment to reflect the reservation costs at the 3PL. If the capacity of the primary transportation option is exceeded on the shipment day, variable costs c_2 arise for $(m - Cap)^+$ demanded units shipped by the alternative option, where m denotes the realized shipment quantity. Note that $m = 0$ also causes fixed costs (reservation at the 3PL); thus, we assume that fixed costs for the primary option are charged for $m \in \mathbb{N}_0$, yielding the following expression for the expected shipment costs per time unit:

$$\begin{aligned} TSC(R,T,Cap) &= \frac{1}{T}\sum_{m=0}^{\infty} Pr(M=m)\Big(\omega(Cap) + c_2(m-Cap)^+\Big) \\ &= \frac{1}{T}\sum_{m=0}^{\infty} Pr(M=m)\Big(c_1 Cap + c_2(m-Cap)^+\Big), \end{aligned} \tag{4.3}$$

where M is defined as a random variable representing the shipment quantity from the warehouse to the production facility on the day of shipment. Hence, $Pr(M = m)$ represents the PMF of the shipment quantity on a shipment day.

To obtain $TSC(R,T,Cap)$, the PMF of M is needed first. For the case without a preorder strategy ($L_d = 0$), we refer to Stenius et al. (2018), whereas $L_d > 0$ and the extension with early deliveries is studied in this chapter. As already mentioned, units that have reached their due date before the shipment date (which are demanded) must be shipped, and units that have not reached the due date (which are ordered) can be shipped in case of available capacity on the primary transportation option, provided that the warehouse has sufficient stock. To describe this shipment policy mathematically, we introduce the following variables:

$M(t)$	Shipment quantity from the warehouse to the production facility at time t
$K(t)$	Remaining units that cannot be shipped at time t due to a lack of stock at the warehouse or limited transportation capacity
$D(s,t)$	Orders at the warehouse during the time interval $(s,t]$, $s < t$
$IL(t)$	Inventory level at time t
$IP(t)$	Inventory position at time t

In the following, we show how the PMF of the shipment quantity can be approximately computed. For this analysis, we distinguish between different cases, which can be separated depending on the length of L_d, L_s and T, and they are shown in Table 4.1. A detailed explanation of all different situations will be provided during the analysis.

Table 4.1: Ranges for different calculations of remaining units

Case	Range of L_d	Range of L_s
1	$L_d \leq T$	$L_s \leq T$
2		$T < L_s \leq T + L_d$
3		$T + L_d < L_s \leq 2T$
4		$L_s > 2T$
5	$T < L_d \leq 2T$	$T < L_s \leq 2T$
6		$2T < L_s \leq T + L_d$
7		$L_s > T + L_d$
8	$2T < L_d \leq 3T$	$2T < L_s \leq T + L_d$
9		$L_s > T + L_d$
10	$L_d > 3T$	$3T < L_s \leq T + L_d$
11		$L_s > T + L_d$

The Cases $L_d \leq T$:

We start our discussion with Case 1 in which $L_d \leq T$ and $L_s \leq T$, also illustrated in Figure 4.2, where t_n $(t_n = nT, n \in \mathbb{N})$ represents the n^{th} shipment day.

The shipment quantity $M(t_n)$ is composed of the following parts. First, the remaining units $K(t_{n-1})$ of the previous shipment day are included in the shipment quantity. These units were backordered due to a lack of stock at the warehouse or ordered units when the due date was not reached, and there was not sufficient capacity available to ship them in advance. Since $L_s \leq T$, all backordered units at time t_{n-1} can be shipped at t_n, and all ordered units before t_{n-1} must at the latest be shipped at t_n. Second, all orders and demands during the interval $(t_{n-1}, t_n]$ can be shipped at t_n if there is sufficient stock and shipment capacity available. Otherwise, backordered or ordered units will be left behind, indicating that we must subtract $K(t_n)$. These considerations lead to the following expression for the shipment quantity at t_n

$$M(t_n) = D(t_{n-1}, t_n) + K(t_{n-1}) - K(t_n). \tag{4.4}$$

This expression reveals that we must characterize the stochastic process describing the remaining units at shipment t_n, denoted by $K(t_n)$. As mentioned above, there are two reasons why units cannot be shipped. First, if there is a lack of stock at the warehouse, then even demanded units cannot be shipped, and $IL(t_n)^-$ units are backordered. Second, all units that are ordered but not demanded must wait until the subsequent shipment

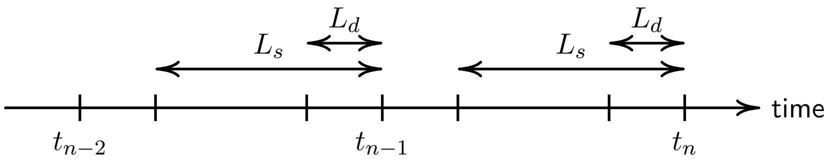

Figure 4.2: Shipment cycle when $L_d < T$ and $L_s \leq T$

dispatches if there is not enough reserved capacity at t_n available. The total number of units that exceeds the capacity of the primary transportation option is given as $(Cap - D(t_{n-1}, t_n) - K(t_{n-1}))^-$, but since all demanded units must be shipped by the alternative transportation option, at maximum, all orders $D(t_n - L_d, t_n)$ remain at the warehouse and must wait until the next shipment. Doing so yields

$$K(t_n) = \max\Big(IL(t_n)^-, \min\Big(D(t_n - L_d, t_n), \Big(Cap - D(t_{n-1}, t_n) - K(t_{n-1})\Big)^-\Big)\Big). \tag{4.5}$$

It is easy to see that the following limit holds,

$$K^\infty(t_n) := \lim_{Cap \to \infty} K(t_n) = IL(t_n)^- \tag{4.6}$$

which shows that, in the case of high capacity, the shipment policy strives to ship all units earlier because only in the case of a lack of stock at the warehouse backorders have to wait until the next dispatch time. In contrast, no early deliveries will occur in the case of small available capacity, as seen in Equation (4.7).

$$K^0(t_n) := \lim_{Cap \to 0} K(t_n) = \max\Big(IL(t_n)^-, D(t_n - L_d, t_n)\Big) \tag{4.7}$$

For the remaining analysis, we split the time interval $(t_{n-1}, t_n]$ into several subintervals, and we express the number of backorders at time t_n based on the inventory position at time $t_n - L_s$, $IP(t_n - L_s)$, because $t_n - L_s$ is the last time when a replenishment order can be placed that arrives before or at t_n. Reformulating Equation (4.5) leads to

$$\begin{aligned} K(t_n) = \max\Big(&\Big(IP(t_n - L_s) - D(t_n - L_s, t_n - L_d) - D(t_n - L_d, t_n)\Big)^-, \\ &\min\Big(D(t_n - L_d, t_n), \Big(Cap - D(t_{n-1}, t_n - L_s) - D(t_n - L_s, t_n - L_d) \\ &- D(t_n - L_d, t_n) - K(t_{n-1})\Big)^-\Big)\Big). \end{aligned} \tag{4.8}$$

Although the distributions of the demand and the inventory position are known and independent, the conditional distribution $Pr(K(t_n) = j|K(t_{n-1}) = i)$ has to be considered; therefore, the PMF of $K(t_n)$ cannot be computed directly by applying Equation (4.8). When computing $K(t_{n-1})$, the number of backorders is obtained based on the inventory position at $t_{n-1} - L_s$, which is why we must also include information about the demand and the inventory position before time t_{n-1}, and we replace $IP(t_n - L_s)$ with

$$IP(t_n - L_s) = \underset{R,Q}{\operatorname{mod}}\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s)\Big). \tag{4.9}$$

We obtain

$$\begin{aligned} K(t_n) = \max\Bigg(\bigg(&\underset{R,Q}{\operatorname{mod}}\Big((IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\ &- D(t_{n-1} - L_d, t_{n-1}) - D(t_{n-1}, t_n - L_s)\Big) - D(t_n - L_s, t_n - L_d) \\ &- D(t_n - L_d, t_n)\bigg)^{-}, \min\Big(D(t_n - L_d, t_n), \big(Cap - D(t_{n-1}, t_n - L_s) \\ &- D(t_n - L_s, t_n - L_d) - D(t_n - L_d, t_n) - K(t_{n-1})\big)^{-}\Big)\Bigg). \end{aligned} \tag{4.10}$$

To obtain $K(t_n)$, we first have to compute $K(t_{n-1})$, which can be specified as

$$\begin{aligned} K(t_{n-1}) = \max\Big(&(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\ &- D(t_{n-1} - L_d, t_{n-1}))^{-}, \min(D(t_{n-1} - L_d, t_{n-1}), (Cap \\ &- D(t_{n-2}, t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\ &- D(t_{n-1} - L_d, t_{n-1}) - K(t_{n-2}))^{-})\Big). \end{aligned} \tag{4.11}$$

This expression means that the number of remaining units $K(t_n)$ depends on the number of remaining units of all previous shipment intervals. We assume that the impact of these quantities decreases the further we look back into the past; therefore, we only include the information of the previous two shipment intervals to determine $K(t_n)$. However, we have to deal with $K(t_{n-2})$. We suggest replacing it with a constant value and providing more information about this approximation at the end of this section.

Hence, the shipment quantity given in Equation (4.4) can be modified, and we replace in Equation (4.12) the expressions for $K(t_n)$ and $K(t_{n-1})$ with Equation (4.10) and

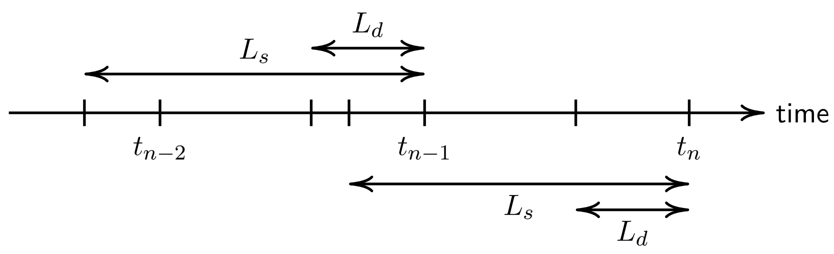

Figure 4.3: Shipment cycle when $L_d \leq T$ and $T < L_s \leq T + L_d$

Equation (4.11) to compute the PMF for the shipment quantity.

$$M(t_n) = D(t_{n-1}, t_n - L_s) + D(t_n - L_s, t_n - L_d) + D(t_n - L_d, t_n) + K(t_{n-1}) - K(t_n) \tag{4.12}$$

As the demand follows a Poisson process, the demand during a specific time interval is Poisson distributed. Further, $IP(t_{n-1} - L_s)$ is uniformly distributed between $R + 1$ and $R + Q$ (Axsäter, 2015). Since all random variables are independently distributed, convolutions can be used to calculate the PMF of the shipment quantity M. For the determination of $K(t_{n-2})$, we rely on an approximation presented at the end of this section. Due to our simplifying assumptions, the obtained distribution is an approximation. Its performance is tested in a numerical study in Section 4.4.

We obtain three additional cases when $L_d \leq T$ and $L_s > T$ based on the length of L_s, as shown in Table 4.1. In Case 2, the supply lead time is in the range of $T < L_s \leq T + L_d$, which is why the time point $t_{n-1} - L_d$ is before $t_n - L_s$, as illustrated in Figure 4.3. The sequence of the latter time points changes if $T + L_d < L_s \leq 2T$, which legitimates Case 3. Case 4 occurs if $L_s > 2T$ because $t_n - L_s$ is before t_{n-2}. The main difference from Case 1 is that the time point $t_n - L_s$ of the last warehouse replenishment order, which will arrive at the latest at t_n, is before the previous shipment day at t_{n-1}, which is why we cannot assure that all backorders at an arbitrary shipment date can be satisfied by the following shipment day. This change in sequences leads to different time intervals during $t_{n-1} - L_s$ and t_n and, therefore, to different formulas of $K(t_n)$ and $K(t_{n-1})$, which are shown in Appendix 4.6.1.

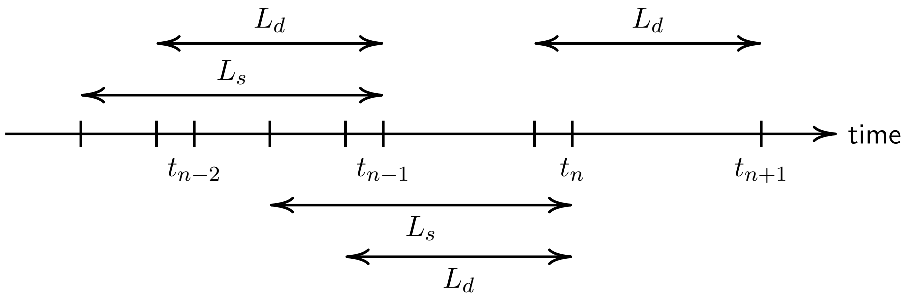

Figure 4.4: Shipment cycle when $T < L_d \leq 2T$ and $T < L_s \leq 2T$

The Cases $L_d > T$:

Now, we focus on all seven cases in which $L_d > T$. To express the differences between $L_d \leq T$ and $L_d > T$, we investigate Case 5 with $T < L_d \leq 2T, T < L_s \leq 2T$, as illustrated in Figure 4.4 in more detail. The other situations can be handled similarly. The shipment policy only allows for shipping an ordered unit one shipment date earlier to the production facility. As long as $L_d \leq T$, this assumption is fulfilled automatically. However, this is not true for $L_d > T$. Orders during $(t_{n+1} - L_d, t_n]$ cannot be shipped at t_n because their official shipment date is at t_{n+2}. The earliest shipment date for these orders is t_{n+1}. Therefore, the demand during $t_{n+1} - L_d$ and t_n is directly added to the number of remaining units at t_n, which can be obtained by

$$\begin{aligned}K(t_n) = D(t_{n+1} - L_d, t_n) + \max\Bigg(\bigg(\underset{R,Q}{\mathrm{mod}}\Big(IP(t_{n-1} - L_s)\\ - D(t_{n-1} - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-2}) - D(t_{n-2}, t_n - L_s)\Big)\\ - D(t_n - L_s, t_n - L_d) - D(t_n - L_d, t_{n-1}) - D(t_{n-1}, t_{n+1} - L_d)\bigg)^-,\\ \min\Big(D(t_n - L_d, t_{n-1}) + D(t_{n-1}, t_{n+1} - L_d), \big(Cap - D(t_{n-1}, t_{n+1} - L_d)\\ - K(t_{n-1})\big)^-\Big)\Bigg). \qquad (4.13)\end{aligned}$$

This property can also be applied for determining $K(t_{n-1})$, where the orders $D(t_n - L_d, t_{n-1})$ will be considered on the shipment day at t_n. Doing so yields

$$\begin{aligned}K(t_{n-1}) = D(t_n - L_d, t_{n-1}) + \max\bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d)\\ - D(t_{n-1} - L_d, t_{n-2}) - D(t_{n-2}, t_n - L_s) - D(t_n - L_s, t_n - L_d)\Big)^-,\\ \min\Big(D(t_{n-1} - L_d, t_{n-2}) + D(t_{n-2}, t_n - L_s) + D(t_n - L_s, t_n - L_d),\\ \big(Cap - K(t_{n-2})\big)^-\Big)\bigg). \qquad (4.14)\end{aligned}$$

The shipment quantity is the sum of demands during the considered shipment interval plus $K(t_{n-1})$ minus $K(t_n)$, as given in Equation (4.4).

$$\begin{aligned} M(t_n) = D(t_{n-1}, t_{n+1} - L_d) + D(t_{n+1} - L_d, t_n) \\ + D(t_n - L_d, t_n) + K(t_{n-1}) - K(t_n) \end{aligned} \tag{4.15}$$

When considering the range $T < L_d \leq 2T$ for the demand lead time, the sequence of time points changes depending on the length of L_s, which is why Case 6 and Case 7 arise. The time point $t_n - L_s$ is between $t_{n-1} - L_d$ and t_{n-2} if $2T < L_s \leq T + L_d$ and accordingly between $t_{n-1} - L_s$ and $t_{n-1} - L_d$ if $T + L_d < L_s$. Similarly, all ranges for the remaining cases can be obtained. All formulas for $K(t_n)$ and $K(t_{n-1})$ can be found in Appendix 4.6.1.

Iterative Procedure to Determine $K(t_{n-2})$

While we have derived expressions for $K(t_n)$ and $K(t_{n-1})$, it remains to determine $K(t_{n-2})$ to be able to compute the PMF of the shipment quantity $M(t_n)$. As mentioned before, we neglect the remaining units of prior shipment days before t_{n-2} and use a fixed value $\bar{K}$ for $K(t_{n-2})$, which is updated in an iterative procedure. The idea is to replace the number of remaining units with its expectation. Therefore, we start with a given value of $\bar{K}$, determine the PMF of $K(t_n)$, and use it to compute the expectation of $K(t_n)$. We also consider a factor that reflects the capacity utilization of the vehicle when we update the value for $\bar{K}$. In general, the obtained value for $\bar{K}$ is not an integer, rendering the computation of the convolutions impossible. As a solution, we determine the PMF for the shipment quantity and the remaining units at t_n two times. First, we use the rounded-down value $\lfloor \bar{K} \rfloor$ and the second time the rounded-up value $\lceil \bar{K} \rceil$. Both results are merged as follows:

$$\begin{aligned} Pr(M(t_n) = m) = \Big(\lceil \bar{K} \rceil - \bar{K}\Big) Pr\Big(M(t_n) = m | \lceil \bar{K} \rceil\Big) \\ + \Big(\bar{K} - \lfloor \bar{K} \rfloor\Big) Pr\Big(M(t_n) = m | \lfloor \bar{K} \rfloor\Big) \end{aligned} \tag{4.16}$$

As a starting value for the iterative procedure, we have chosen $(Cap - \lambda T - \frac{1}{2}\lambda L_d)^-$, which is obtained by replacing the random demand in the expression with the total number of units, which exceeds the capacity of the primary transportation option, with its expectation. Furthermore, the remaining units $K(t_{n-1})$ are replaced with 50 % of the expected number of ordered units with a due date after the shipment time. In summary, we provide a sketch of the algorithm to compute the expected shipment costs.

Step 1: Start with $\bar{K} = \left(Cap - \lambda T - \frac{1}{2}\lambda L_d\right)^-$

Step 2: Compute for $\lceil \bar{K} \rceil$ the PMF of $K(t_n)$ from Equation (4.10) and Equation (4.11) for Case 1, or use the corresponding formulas for the other cases

Step 3: Compute for $\lfloor \bar{K} \rfloor$ the PMF of $K(t_n)$ from Equation (4.10) and Equation (4.11) for Case 1, or use the corresponding formulas for the other cases

Step 4: Compute the PMF of $K(t_n)$, similar to Equation (4.16)

Step 5: Compute the expected number of remaining units at t_n by $\bar{K}(t_n) = \sum_{j=0}^{\infty} jPr(K(t_n) = j)$

Step 6: If $|\bar{K} - \bar{K}(t_n)\frac{Cap}{\lambda T}| < 0.1$, go to step 7; otherwise set $\bar{K} = \left\lfloor \bar{K}(t_n)\frac{Cap}{\lambda T} \cdot 10 \right\rceil : 10$ and go to step 2

Step 7: Compute the PMF of the shipment quantity according to Equation (4.16)

Step 8: Compute $TSC(R, T, Cap)$ with current PMF of the shipment quantity by Equation (4.3).

4.2.2 Expected Inventory Costs at the Warehouse

For the analysis of the expected inventory costs per time unit, we adapt the methodology introduced in Marklund (2011), who studied an inventory system with time-based shipment consolidation without ADI and without flexible deliveries, reflecting the Case $L_d = 0$. Therefore, we only discuss the Case $L_d > 0$ with flexible deliveries in this chapter. We observe each unit going through the system separately, calculate the cost for each unit, and then consider the expectation of this cost. In the first step, we show how the (R, Q) policy is connected to a base-stock policy with base-stock level S because we use this relationship in our further analysis.

Let us denote by t_r the time when the warehouse just placed a replenishment order of size Q at the outside supplier, which arrives after a supply lead time L_s. The units of this order are then consumed in the defined order $1, 2, \ldots, Q$. The first unit of this batch is needed for the $(R+1)^{th}$ order at the warehouse after t_r because, at t_r, there are still R units on stock that are used first. The observation of the first unit represents the situation in which the warehouse uses an $(S-1, S)$ policy with base-stock level $S = R+1$. Similarly, we can relate a base-stock policy to each unit of the batch, for example, using a base-stock level of $S = R+2$ for the second unit and finally $S = R+Q$ for the final unit of the batch. Thus, one possibility for obtaining the system's expected total inventory costs is to replace the (R, Q) policy with Q base-stock policies with base-stock levels $S = R+1, R+2, \ldots, R+Q$. Therefore, the expected inventory costs of the system using an (R, Q) policy can be calculated as shown in Equation (4.17), where $TIC(S, T, Cap)$ represents the expected inventory cost of the system per time unit when using a base-stock policy with base-stock level S and a shipment interval T with a capacity reservation of Cap units.

$$TIC(R, T, Cap) = \frac{1}{Q} \sum_{S=R+1}^{R+Q} TIC(S, T, Cap) \tag{4.17}$$

In Case $S > 0$, the warehouse orders the *considered unit* at the outside supplier before a facility orders it, whereas Case $S \leq 0$ implies that the warehouse orders the considered unit at the outside supplier after or at the same time at which the facility orders it at the warehouse. These two cases have to be discussed separately, and we have to introduce some additional notation for further analysis:

$\Omega(x)$ Length of the time interval between the replenishment moment at the warehouse and the moment when the x^{th} order from the production facility arrives at the warehouse, random variable

V Shipment delay, defined as the length of the time interval between the time when a unit is demanded and available for a shipment and the subsequent shipment day, random variable

$G^x(t)$ Cumulative distribution function of an Erlang (x, λ) distribution

$g^x(t)$ Probability density function of an Erlang (x, λ) distribution

$U(t)$ Cumulative distribution function of a uniformly distributed random variable on the interval $[0, T]$

$u(t)$ Probability density function of a uniformly distributed random variable on the interval $[0, T]$

We call a unit *prequalified* if it is available and ordered because it can be shipped before it is actually demanded since we allow for early deliveries. A unit is *qualified* if it is available and demanded and thus has to be shipped with the primary or alternative transportation option on the subsequent shipment day.

The Case $S > 0$:

Observing a specific unit on its way through the system, we recognize four essential time instances. First, a unit is available for satisfying demand at the warehouse exactly L_s time units after the replenishment order was placed, thus at time $t_a = t_r + L_s$. The next important event is the time when the facility order for this unit arrives at the warehouse, which is given as $t_o = t_r + \Omega(S)$, where $\Omega(S)$ represents the time until the S^{th} facility order occurs. The third interesting time is when the status of the unit changes from an order to a demand, which occurs at $t_d = t_r + \Omega(S) + L_d$. Finally, there is the time t_s when the unit is shipped to the production facility, depending on the available reserved transportation capacity.

Which types of costs occur and to what extent depend on the sequence of the events. For the derivation of the formulas, we have to distinguish all possible Situations $i \in \{A, B, C, D, E, F, G\}$. Situation A is illustrated in Figure 4.5, and we explain all of the situations in the following in more detail. The associated figures are illustrated in Appendix 4.6.2.

The first three situations (Situation A, Situation B, Situation C) are related to a situation in which the unit is available when the order occurs, which means that $t_a < t_o$. Otherwise, the unit is ordered while there is no available stock for the considered unit at the warehouse (Situation D, Situation E, Situation F, Situation G). We further differentiate between the situations in which prequalification and qualification occur at different points in time (Situation D, Situation E, Situation F) and at the same time (Situation G), whereby

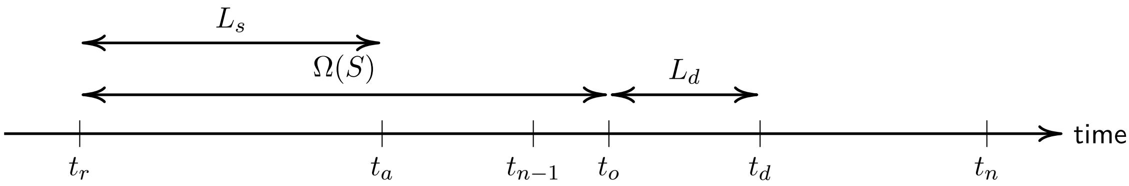

Figure 4.5: Important moments in time in Situation A

the latter situation only occurs if the unit is available after it is demanded. If the unit is available before demand occurs, then the unit is prequalified at time $\max(t_a, t_o)$ and qualified at time t_d. In the following, $t_n \in \mathbb{N}$ denotes an arbitrary shipment day.

A $t_a < t_o, t_{n-1} < t_o < t_d < t_n, t_s = t_n$
Since prequalification and qualification occur in the same shipment interval, the unit must be shipped on the next shipment day ($t_s = t_n$), indicating there are no early-delivery costs, and the unit must wait V time units. Holding costs are charged for a time interval with length $\Omega(S) - L_s + L_d + V$.

B $t_a < t_o, t_o < t_{n-1} < t_d < t_n, t_s = t_{n-1}$
This situation can only occur if there is sufficient reserved transportation capacity available at t_{n-1} to initiate an early delivery, which is then associated with early-delivery costs for $T - V$ time units. The unit is kept on stock for $\Omega(S) - T - L_s + L_d + V$ time units.

C $t_a < t_o, t_o < t_{n-1} < t_d < t_n, t_s = t_n$
There is not sufficient transportation capacity available to allow for early delivery of the considered unit in this situation. Then, late-delivery costs for V time units are due, and holding costs are incurred for $\Omega(S) - L_s + L_d + V$ time units.

D $t_o \leq t_a, t_{n-1} < t_a < t_d < t_n, t_s = t_n$
Similar to Situation A, there are no early-delivery costs, and late-delivery costs are incurred for V time units. The unit is kept on stock for $\Omega(S) - L_s + L_d + V$ time units, which is shorter than A because the unit is available after t_{n-1}.

E $t_o \leq t_a, t_a < t_{n-1} < t_d < t_n, t_s = t_{n-1}$
This situation can only occur if there is sufficient reserved transportation capacity available at t_{n-1} to initiate an early delivery, which is then associated with early-delivery costs for $T - V$ time units. The unit is kept on stock for $\Omega(S) - T - L_s + L_d + V$ time units.

Table 4.2: Holding, late-delivery, and early-delivery costs expressions for all situations

Situation i	Holding costs	Late-delivery costs	Early-delivery costs
A	$h(\Omega(S) - L_s + L_d + V)$	lV	0
B	$h(\Omega(S) - T - L_s + L_d + V)$	0	$e(T - V)$
C	$h(\Omega(S) - L_s + L_d + V)$	lV	0
D	$h(\Omega(S) - L_s + L_d + V)$	lV	0
E	$h(\Omega(S) - T - L_s + L_d + V)$	0	$e(T - V)$
F	$h(\Omega(S) - L_s + L_d + V)$	lV	0
G	hV	$l(V + L_s - L_d - \Omega(S))$	0

F $t_o \leq t_a, t_a < t_{n-1} < t_d < t_n, t_s = t_n$
This situation is comparable to C such that late-delivery costs for V time units are, due and holding costs are incurred for $\Omega(S) - L_s + L_d + V$ time units.

G $t_o \leq t_a, t_o < t_d < t_a < t_n, t_s = t_n$
Since prequalification and qualification occur at the same moment in time, there are no early-delivery costs. Inventory holding costs only occur for V time units, and late-delivery costs (including the backorder costs) are charged for $V + L_s - L_d - \Omega(S)$ time units.

In summary, we provide all of the cost expressions for situations $i \in \{A, B, C, D, E, F, G\}$ in Table 4.2.

Obviously, for a given S, T, and Cap, the inventory costs depend on the random variables $\Omega(S)$ and V and the available reserved transportation capacity at t_n. However, we only know the PMF of the shipment quantity at t_n after all demands and orders of the considered shipment interval occurred, which again is why we rely on an approximation. We denote the inventory costs for Case $S > 0$ in the following analysis by $C(\Omega(S), V, M(t_n))$. Denoting the joint density function of $\Omega(S)$, V and $M(t_n)$ by $f(x, y, z)$, we can obtain the expected inventory costs by

$$\begin{aligned} E[C(\Omega(S), V, M(t_n))] &= \int_0^\infty \int_0^T \int_0^\infty C(x, y, z) f(x, y, z) dz\, dy\, dx \\ &= \sum_{\substack{i \in \{A,B,C,\\ D,E,F,G\}}} \int_0^\infty \int_0^T \int_0^\infty C_i(x, y, z) f_i(x, y, z)\, dz\, dy\, dx, \end{aligned} \tag{4.18}$$

which means that we can split the derivation into several parts, where each part corresponds to one of the aforementioned situations. For the following analysis, we neglect the dependency between $M(t_n)$ and the other random variables such that we obtain

$$E[C(\Omega(S), V, M(t_n))] = \sum_{\substack{i\in\{A,B,C,\\ D,E,F,G\}}} \int_0^\infty \int_0^T \int_0^\infty C_i(x,y,z) f_i(x,y) f(z) dz\, dy\, dx. \quad (4.19)$$

Therefore, we can reformulate Equation (4.19) as

$$\begin{aligned} &E[C(\Omega(S), V, M(t_n))] \\ &\quad = \sum_{\substack{i\in\{A,B,C,\\ D,E,F,G\}}} \int_0^\infty \int_0^T \int_0^\infty C_i(x,y) f_i(x,y) f(z) dz\, dy\, dx \\ &\quad = \sum_{i\in\{A,D,G\}} \int_0^\infty \int_0^T C_i(x,y) f_i(x,y)\, dy\, dx \\ &\quad\quad + Pr(M(t_n) < Cap) \sum_{i\in\{B,E\}} \int_0^\infty \int_0^T C_i(x,y) f_i(x,y)\, dy\, dx \\ &\quad\quad + Pr(M(t_n) \geq Cap) \sum_{i\in\{C,F\}} \int_0^\infty \int_0^T C_i(x,y) f_i(x,y)\, dy\, dx, \end{aligned} \quad (4.20)$$

where $C_i(x,y)$ and $f_i(x,y)$ represent the inventory costs depending on $\Omega(S)$ and V and the joint probability density function of $\Omega(S)$ and V for Situation $i \in \{A,B,C,D,E,F,G\}$, respectively.

Since the cost expressions can easily be derived from Table 4.2, it remains to determine the functions $f_i(x,y), i \in \{A,B,C,D,E,F,G\}$. It can be shown that the functions $f_i(x,y)$, are positive on different domains (see Table 4.3). Further, in the range where they are unequal to zero, the functions have the form $g^S(x)u(y)$ for all Situations $i \in \{A,B,C,D,E,F,G\}$. In the following, we show the derivation for Situation A, whereas the remaining derivations are given in Appendix 4.6.2.

Table 4.3: Range for x and y dependent on Situation $i \in \{A,B,C,D,E,F,G\}$

Situation i	x	y
A	$L_s < x < \infty$	$0 \leq y \leq (T-L_d)^+$
B	$L_s < x < \infty$	$(T-L_d)^+ < y \leq T$
C	$L_s < x < \infty$	$(T-L_d)^+ < y \leq T$
D	$L_s - L_d < x \leq L_s$	$0 \leq y \leq (T-L_d+L_s-x)^+$
E	$L_s - L_d < x \leq L_s$	$(T-L_d+L_s-x)^+ < y \leq T$
F	$L_s - L_d < x \leq L_s$	$(T-L_d+L_s-x)^+ < y \leq T$
G	$0 \leq x \leq L_s - L_d$	$0 \leq y \leq T$

The joint cumulative distribution function of $\Omega(S)$ and V is denoted by $F_i(x,y)$ for Situation $i \in \{A,B,C,D,E,F,G\}$ and for Situation A given by

$$\begin{aligned}
F_A(x,y) &= Pr\Big(\Omega(S) \leq x, V \leq y, t_a < t_o, t_{n-1} < t_o < t_d < t_n, t_s = t_n\Big) \\
&= Pr\Big(\Omega(S) \leq x, V \leq y, t_r + L_s < t_r + \Omega(S), \\
&\quad t_{n-1} < t_r + \Omega(S) < t_r + \Omega(S) + L_d < t_n\Big) \\
&= Pr\Big(\Omega(S) \leq x, V \leq y, L_s < \Omega(S), \\
&\quad 0 < t_r + \Omega(S) - t_{n-1} < t_r + \Omega(S) + L_d - t_{n-1} < t_n - t_{n-1}\Big) \\
&= Pr\Big(\Omega(S) \leq x, V \leq y, L_s < \Omega(S), \\
&\quad 0 < t_r + \Omega(S) - t_{n-1} < t_r + \Omega(S) + L_d - t_{n-1} < T\Big). \qquad (4.21)
\end{aligned}$$

Since $T = V + L_d + t_r + \Omega(S) - t_{n-1}$ (see Figure 4.5), we obtain for $x > L_s$ and $y \leq (T - L_d)^+$.

$$\begin{aligned}
F_A(x,y) &= Pr\Big(L_s < \Omega(S) \leq x, V \leq y, \\
&\quad 0 < t_r + \Omega(S) - t_{n-1} < t_r + \Omega(S) + L_d - t_{n-1} < T\Big) \\
&= Pr\Big(L_s < \Omega(S) \leq x, V \leq y, 0 < T - V - L_d < T - V < T\Big) \\
&= Pr\Big(L_s < \Omega(S) \leq x, V \leq y, V < (T - L_d)^+\Big) \\
&= Pr\Big(L_s < \Omega(S) \leq x, V \leq y\Big) \\
&= \Big(G^S(x) - G^S(L_s)\Big)U(y) \qquad (4.22)
\end{aligned}$$

V is uniformly distributed between $[0,T]$ (Tijms, 2003), which is why the upper bound of y cannot be smaller than 0. Thus, $f_A(x,y)$ is given as the partial derivative with respect to x and y

$$f_A(x,y) = \begin{cases} g^S(x)u(y) & L_s < x < \infty, 0 \leq y \leq (T - L_d)^+ \\ 0 & \text{otherwise.} \end{cases} \qquad (4.23)$$

Since demand follows a Poisson process, $\Omega(S)$ is Erlang distributed with the parameters S and λ and is independent of the shipment delay V (Tijms, 2003). Therefore, the expectation of the inventory costs for each situation can be computed, and the results for

$$E[C_i(\Omega(S),V)] = \int_0^\infty \int_0^T C_i(x,y) f_i(x,y)\, dy\, dx \quad \forall\, i \in \{A,B,C,D,E,F,G\} \tag{4.24}$$

are given in Table 4.4 and Table 4.5. Since the range for y depends on T and L_d, we must consider Case $L_d \leq T$ and Case $L_d > T$ separately. The derivation is again reported in Appendix 4.6.2.

The Case $S \leq 0$:

If $S = 0$, then the considered unit is ordered by the production facility at the same time that the warehouse orders it from the outside supplier. Thus, no safety stock is kept at the warehouse. Focusing on $S < 0$, the warehouse orders the considered unit at the outside supplier after the next $|S|^{th}$ facility orders occur. Due to $L_d \leq L_s$, both cases imply that the considered unit is always demanded before it is available for shipment. Therefore, warehouse holding costs arise for the time interval V, whereas late-delivery costs occur for the time interval $V + L_s - L_d + \Omega(|S|)$. As $t_a \geq t_d$, the flexible delivery option cannot be used. Therefore, inventory costs $\tilde{C}(\Omega(S),V)$ for Case $S \leq 0$ are independent of $M(t_n)$, and the expectation of this cost can be obtained by

$$E[\tilde{C}(\Omega(S),V)] = h\frac{T}{2} + l\left(\frac{|S|}{\lambda} + L_s - L_d + \frac{T}{2}\right). \tag{4.25}$$

In summary, the expected inventory costs of the system are given by

$$TIC(S,T,Cap) = \begin{cases} \lambda E\Big[C(\Omega(S),V,M(t_n))\Big], & \text{for } S > 0 \\ \lambda E\Big[\tilde{C}(\Omega(S),V)\Big], & \text{for } S \leq 0. \end{cases} \tag{4.26}$$

4.3 Approximation Method

This section shows how to minimize the expected total cost $TC(R,T,Cap)$ per time unit when $L_d \geq 0$ by determining the near-optimal reorder level R^*_{app}, the near-optimal shipment interval T^*_{app} as well as the near-optimal capacity reservation Cap^*_{app}.

Table 4.4: Expected inventory costs conditioned on the different situations for the Case $L_d \leq T$

i	$E[C_i(\Omega(S), V)]$
A	$h\frac{T-L_d}{T}\Big(\Big(\frac{T+L_d}{2} - L_s\Big)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Big) + l\frac{(T-L_d)^2}{2T}\Big(1 - G^S(L_s)\Big)$
B	$h\frac{L_d}{T}\Big(\frac{L_d}{2} - L_s\Big)\Big(1 - G^S(L_s)\Big)\frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Big) + e\frac{L_d^2}{2T}\Big(1 - G^S(L_s)\Big)$
C	$h\frac{L_d}{T}\Big(\frac{L_d}{2} - L_s + T\Big)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Big) + l\frac{2TL_d - L_d^2}{2T}\Big(1 - G^S(L_s)\Big)$
D	$h\Big(\frac{T^2-(L_s-L_d)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) + \frac{(L_s-L_d)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) - \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$ $+l\Big(\frac{(T-L_d+L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) - \frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) + \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$
E	$h\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big)) + \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$ $+e\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) + \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$
F	$h\Big(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s+T)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) + \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$ $+l\Big(\frac{2T(L_d-L_s)-(L_d-L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) + \frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) - \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big)$
G	$h\frac{T}{2}G^S(L_s - L_d) + l\Big(\Big(L_s - L_d + \frac{T}{2}\Big)G^S(L_s - L_d) - \frac{S}{\lambda}G^{S+1}(L_s - L_d)\Big)$

Table 4.5: Expected inventory costs conditioned on the different situations for the Case $L_d > T$

i	$E[C_i(\Omega(S),V)]$
A	0
B	$h\Big(\Big(L_d - L_s - \frac{T}{2}\Big)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Big) + e\frac{T}{2}\Big(1 - G^S(L_s)\Big)$
C	$h\Big(\Big(L_d - L_s + \frac{T}{2}\Big)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Big) + l\frac{T}{2}\Big(1 - G^S(L_s)\Big)$
D	$h\Big(\frac{T^2-(L_s-L_d)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) + \frac{(L_s-L_d)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big)\Big)$ $+l\Big(\frac{(T-L_d+L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) - \frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big)\Big)$
E	$h\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) + \Big(L_d - L_s - \frac{T}{2}\Big)\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big) + \frac{S}{\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d + T)\Big)\Big)$ $+e\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) + \frac{T}{2}\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big)\Big)$
F	$h\Big(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) + \frac{(L_d-L_s+T)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) + \Big(L_d - L_s + \frac{T}{2}\Big)\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big) + \frac{S}{\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d + T)\Big)\Big)$ $+l\Big(\frac{2T(L_d-L_s)-(L_d-L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) + \frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big)$ $-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) + \frac{T}{2}\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big)\Big)$
G	$h\frac{T}{2}G^S(L_s - L_d) + l\Big(\Big(L_s - L_d + \frac{T}{2}\Big)G^S(L_s - L_d) - \frac{S}{\lambda}G^{S+1}(L_s - L_d)\Big)$

Examples reveal that $TC(R,T,Cap)$ is not jointly convex in R, T, and Cap. However, during all of our numerical experiments, we could not find any example in which $TC(R,T,Cap)$ was not convex in T for a fixed R and Cap (not convex in Cap for a fixed R and T). Using this property, we perform a bounded enumeration. In our numerical study, we focus on the optimization of R^*_{app} and T^*_{app} for a fixed Cap and on the optimization of R^*_{app} and Cap^*_{app} for a fixed T, respectively. T and Cap are highly dependent on each other since high utilization of the reserved transportation capacity mainly depends on the expected number of orders during a shipment interval, thus on λT. In general, a modified given Cap changes the optimal T^*_{app}, and inversely, a modified given T changes the optimal Cap^*_{app}. To limit the computational time, we assume that either Cap or T is given. This is also sufficient to answer the research questions.

First, we define a lower and an upper bound on the optimal decision variable R^*_{app}, denoted by R^l and R^u, respectively. A proven lower bound is $R^l = -Q$ (Axsäter, 1998). When increasing R, there is a point at which another increase in R does not influence the waiting, early-delivery, and shipment costs because the stock on hand remains sufficient to always satisfy all orders and demands; thus, the flexible delivery option is already used to some extent. The only effect then is an increase in inventory holding costs. The reorder level is high enough when the demand during the supply lead time never exceeds R, indicating that backorders do not occur. Therefore, $R^u = \min(R : Pr(D(0, L_s) > R) < \varepsilon)$, where ε is a small number close to zero. Note that the bounds do not depend on T or on Cap.

4.3.1 Determination of R^*_{app} and T^*_{app} for a given Cap

Since the shipment interval can only take natural numbers, we define the lower bound on T^*_{app} as $T^l = 1$. The upper bound on T^*_{app} is $T^u(R, Cap)$ and can be found for each R and a given Cap individually, where we use the convexity property. For a given $R \in \{R^l, R^l + 1, \ldots, R^u\}$ and Cap, we increase T by 1 until $TC(R, T-1, Cap) < TC(R, T, Cap)$.

Step 1: Determine $R^u = \min(R : Pr(D(0, L_s) > R) < \varepsilon)$ and $R^l = -Q$, and fix Cap

Step 2: For all given $R \in \{R^l, R^l + 1, \ldots, R^u\}$, compute $TC(R, T, Cap)$ with $T = T^l = 1$

Step 3: Increase T by 1 and compute $TC(R, T, Cap)$ for all relevant R

Step 4: If $TC(R, T-1, Cap) < TC(R, T, Cap)$ for all $R = R^l, R^l+1, \ldots, R^u$, go to step 5; else continue with step 3

Step 5: Find R^*_{app} and T^*_{app}, which minimize the expected total cost $TC(R, T, Cap)$ for a fixed Cap.

4.3.2 Determination of R^*_{app} and Cap^*_{app} for a given T

Obviously, the transportation capacity reserved cannot be negative and has to be integer. Therefore, the lower bound on Cap^*_{app} is $Cap^l = 0$, where only the alternative transportation option can be used why early deliveries are not allowed at all. The upper bound on Cap^*_{app} is $Cap^u(R, T)$ and can be found for each R and a given T individually, where we use the convexity property. For a given $R \in \{R^l, R^l+1, \ldots, R^u\}$ and T, we increase Cap by 1 until $TC(R, T, Cap-1) < TC(R, T, Cap)$.

Step 1: Determine $R^u = \min(R : Pr(D(0, L_s) > R) < \varepsilon)$ and $R^l = -Q$, and fix T

Step 2: For all given $R \in \{R^l, R^l+1, \ldots, R^u\}$, compute $TC(R, T, Cap)$ with $Cap = Cap^l = 0$

Step 3: Increase Cap by 1 and compute $TC(R, T, Cap)$ for all relevant R

Step 4: If $TC(R, T, Cap-1) < TC(R, T, Cap)$ for all $R = R^l, R^l+1, \ldots, R^u$, go to step 5; else continue with step 3

Step 5: Find R^*_{app} and Cap^*_{app}, which minimize the expected total cost $TC(R, T, Cap)$ for a fixed T.

4.4 Numerical Study

In this section, we first present the results of a numerical study to investigate the performance of the applied approximation for the expected total cost $TC(R, T, Cap)$. Second, we investigate the influences of ADI and the flexible shipment consolidation program on the expected total cost, as well as on the variables to be optimized.

In the numerical study, we focus on bulky and expensive items to show how a company can apply the presented model. We start with a definition of a base case, where the parameters related to the inventory system are in a similar range as in Marklund (2011) and other references. The order rate of the item at the warehouse is given as $\lambda = 2$. The holding costs parameter h at the warehouse equals 1 per unit and time unit, whereas the late-delivery and early-delivery costs parameters at the warehouse are fixed to $l = 2$ and $e = 2$ per unit and time unit, respectively. Due to a time-based shipment consolidation program, the production facility must hold more safety stock, wherefore we consider late-delivery costs. Early-delivery costs represent holding costs at the production facility. Both reasons justify costs parameters l and e close to h. We rely on similar ranges for late-delivery costs, as, e.g., in Çetinkaya et al. (2008). The order quantity Q equals 10 to limit computational time, and the transportation capacity is fixed at $Cap = 10$, which is reasonable for a bulky product. Shipment costs depend on the fixed costs parameter $\omega(Cap)$ for reserved transportation capacity and the variable costs parameter c_2 in case the reserved capacity is exceeded. In our base case, we fix the variable reservation cost to $c_1 = 10$ and obtain in the base case for $Cap = 10$ a value for the fixed shipment costs $\omega(Cap) = 100$. Additionally, variable shipment costs depend on c_1 according to $c_2 = 2c_1$. When the primary transportation capacity is fully utilized, a cost of c_1 arises per unit shipped for the primary transportation option. We double the unit shipment costs for the alternative transportation option for the base case. A reasonable supply lead time from the outside supplier to the warehouse is $L_s = 2$, whereas $L_d = 1$ time units.

4.4.1 Performance of the Approximation

Before deriving managerial insights, we validate our approximation method with a simulation study. For this study, we focus on the optimization of R and T for a given Cap. Therefore, we define a mixed-level fractional factorial design, which relates to the base case. For parameters λ, Cap, $\frac{w}{h}$ and $\frac{e}{h}$, the base case defines the medium level, which is extended by low and high levels. Since h, l, and e represent inventory holding costs at the warehouse and indirectly at the production facility, we only investigate a changed relation between h and l (between h and e). For parameters c_2, L_d and L_s, we define low and high levels as follows: $c_2 = xc_1$ with $x \in \{1.5, 2\}$, $L_d \in \{1, 2\}$ and $L_s \in \{2, 4\}$. We do not vary c_1 since the relation of cost for reservations and spontaneous shipments changes by modifying c_2. Additionally, we fix the order quantity Q as mentioned in the base case because the replenishment costs do not have an influence on the optimal decisions. This yields $3^4 \cdot 2^3 = 648$ instances.

For these 648 instances, we determine the parameters R^*_{app} and T^*_{app}, solving the optimization problem in Equation (4.2), relying on the results of Section 4.2 and Section 4.3. We use a simulation to evaluate the system's expected total cost of a policy and call this cost the exact expected total cost to validate these results. The length of each simulation run is 52000 days, while the last 50000 days are used for the cost computation. We use sequential sampling and stop if the half-width of the 95 % confidence interval of the average total cost is smaller than 0.5 % of the average total cost of the considered instance.

To determine the optimal policy parameters (R^*, T^*), we combine the simulation with a neighborhood search and use R^*_{app} and T^*_{app} as initial values. The neighborhood includes all points $(R^*_{app} + g, T^*_{app} + G)$ with $g, G \in \{-1, 0, 1\}$. If the neighborhood offers a better average total cost value than the initials, the neighborhood search is repeated for the best value in the neighborhood. This iterative procedure can be stopped if no better average total cost value can be found. The obtained policy parameters are locally optimal and define the optimal decision (R^*, T^*).

We are interested in the relative average total cost increase caused by not making the optimal decision with our approach. Therefore, we calculate the relative total cost difference between the average total cost of the optimal policy $TC_{sim}(R^*, T^*)$ and the average total cost of the policy determined by our approach $TC_{sim}(R^*_{app}, T^*_{app})$ for all instances. $TC_{sim}(R, T)$ represents the average total cost computed by simulation, which we assume to be the correct average total cost value. We define the relative cost difference as

$$\delta TC_{sim} = \frac{TC_{sim}(R^*_{app}, T^*_{app}) - TC_{sim}(R^*, T^*)}{TC_{sim}(R^*, T^*)} 100 \tag{4.27}$$

and also compute aggregate values. In Table 4.6, the aggregated results of all 648 examples are provided to investigate the impact of the input parameters on the performance. The average (maximum) total cost deviation is 0.20 % (10.54 %). The worst case is observed in a situation in which the reserved transportation capacity is small compared to the average demand. In these situations, $K(t_{n-2})$ has a more significant influence on the shipment quantity at t_n, explaining the decreasing performance. However, from an economic point of view, a small reserved capacity is only acceptable if the demand is comparatively low.

For 598 of 648 instances, we found the optimal policy parameters using our approximate

Table 4.6: Average and maximum of the relative cost deviation for all examples

Parameters		Value	Average relative cost deviation in %	Maximum relative cost deviation in %
Late-delivery costs	l	1	0.12	10.28
		2	0.10	5.34
		5	0.38	10.54
Early-delivery costs	e	1	0.24	10.54
		2	0.22	10.28
		5	0.14	7.63
Shipment costs	c_2	$1.5c_1$	0.18	10.54
		$2c_1$	0.22	10.28
Demand rate	λ	1	0.00	0.42
		2	0.04	1.32
		4	0.56	10.54
Demand lead time	L_d	1	0.00	0.00
		2	0.40	10.54
Supply lead time	L_s	2	0.27	10.54
		4	0.13	10.35
Capacity	Cap	5	0.58	10.54
		10	0.01	0.98
		20	0.00	0.34
Total			0.20	10.54

approach. Only in 7.72 % of the instances could we not find the optimal values; however, for 31 of these 50 examples, the optimal values were located in the direct neighborhood $(R^*_{app} + g, T^*_{app} + G)$ with $g, G \in \{-1, 0, 1\}$. For only 19 examples, larger deviations in the optimal policy parameters were observed. The maximum deviation between R^*_{app} and R^* for all 648 examples is four, whereas the deviation between T^*_{app} and T^* is one at maximum.

Focusing only on the 50 examples in which we do not derive the optimal policy parameter, we observe an average (a maximum) total cost increase of 2.58 % (10.54 %). A more detailed look at the results also reveals that a non-optimal shipment interval has a larger effect on the expected total cost than a non-optimal reorder level. Choosing a shipment interval that is one day shorter (longer) than the optimal shipment interval means that more demands must be shipped by the alternative transportation option (the transportation capacity is less utilized on average), explaining this observation.

In summary, we conclude that our approximation has an excellent performance for most of the relevant cases, and even in the other situations, it is acceptable. Therefore, we can use our model to generate managerial insights.

4.4.2 Managerial Insights

In this section, we quantify the added value of ADI under flexible shipment consolidation and investigate the impact of the demand lead time on the optimal shipment interval and on the optimal capacity of the primary transportation option. Therefore, we use a different experimental design to reduce computational time without losing insights and reduce the level for the order rate ($\lambda \in \{1, 2\}$) while we increase the levels for the demand lead time $L_d \in \{0, 2, 4, 6, 8\}$ and fix L_s to 10. The other parameters are the same as in the base case and the former study.

Cost Improvements by Advance Demand Information and Flexible Deliveries

First, we focus on a given capacity and on the influence of an increasing demand lead time on the expected total cost. We compute the average marginal relative decrease in the expected total cost when the demand lead time L_d is stepwise increased by two time units while the other parameters are fixed. The results are presented in Table 4.7 and indicate that large cost reductions can be achieved if customers are willing to place orders in advance. In general, it can be seen that the longer the demand lead time, the greater the total cost reductions. The maximal marginal relative cost reduction when L_d is increased from 0 to 2 is 20.29 %, whereas we can achieve a maximum decrease of the expected total cost when L_d is increased from 0 to 8 of 35.57 %. Further, it can be observed that an increase in the demand lead time of two time units can have a different effect depending on the starting point. For example, coming from the situation where $L_d = 0$ to a situation where $L_d = 2$, on average, the expected total cost per time unit can be decreased from 62.01 to 55.54 (10.43 %), whereas the expected total cost per time unit can be reduced from 55.54 to 52.91 (4.74 %) when increasing the demand lead time from 2 to 4. Thus, the marginal value of the ADI decreases with an increasing amount of information.

Table 4.7: Marginal relative cost reduction enabled by ADI and flexible deliveries

Parameter	Value	Marginal relative cost difference in % $L_d = 0 \to L_d = 2$	Marginal relative cost difference in % $L_d = 2 \to L_d = 4$	Marginal relative cost difference in % $L_d = 4 \to L_d = 6$	Marginal relative cost difference in % $L_d = 6 \to L_d = 8$	Relative cost difference in % $L_d = 0 \to L_d = 8$
l	1	6.11	2.27	0.71	0.32	9.19
	2	8.42	3.88	1.58	0.64	13.91
	5	15.01	7.36	3.57	1.63	25.32
e	1	11.17	6.09	3.26	1.66	20.63
	2	10.73	5.20	2.31	0.9	18.13
	5	9.40	2.96	0.67	0.13	12.79
c_2	$1.5c_1$	9.93	4.73	2.02	0.92	16.70
	$2c_1$	10.91	4.74	2.10	0.89	17.66
λ	1	9.53	5.85	3.11	1.57	18.77
	2	11.00	4.01	1.39	0.49	16.17
Cap	5	11.36	2.66	0.96	0.34	14.83
	10	10.90	4.35	1.25	0.45	16.22
	20	9.33	6.61	3.58	1.74	19.78
Total		10.43	4.74	2.06	0.90	17.18

There are two possible sources for the cost reduction. First, as already observed in (Hariharan and Zipkin, 1995), a longer demand lead time reduces the effective supply lead time $(L_s - L_d)$ and, therefore, the safety stock. Second, due to the flexible shipment policy, a longer demand lead time results in more orders available to be shipped in advance. Thus, better utilization of transportation capacity can be achieved, reducing the need for emergency deliveries. Since early deliveries are only allowed when enough remaining transportation capacity is available, it is clear that the reserved capacity must have an impact on the benefit of ADI in our setting. It can be observed that the marginal relative cost reduction decreases more than average with increasing L_d in instances where $Cap = 5$ (see Table 4.7). The mentioned effect is below-average in instances where $Cap = 20$. Although the marginal relative cost reduction also decreases as L_d increases, it decreases much less, and even when the demand lead time is increased from 2 to 4, 6.61 % of the expected total cost can still be saved. With a large reserved transportation capacity, there is a higher probability of unused capacity for orders being available. In situations with a small reserved transportation capacity, the capacity is already fully utilized on many shipping days when $L_d = 2$, so further information will not result in additional early deliveries in many cases. Since the reserved capacity seems to be an important variable, we will also determine the optimal reserved capacity later.

To investigate the proportion of the average total cost decrease caused by reducing the safety stock at the warehouse and by early shipments, respectively, we computed the

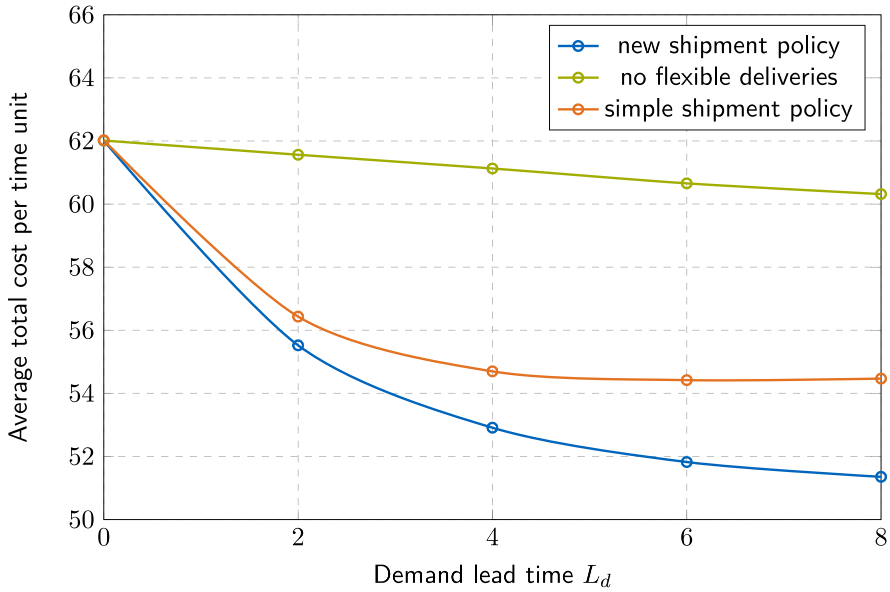

Figure 4.6: Average total cost per time unit across all instances

expected total cost for the same 540 examples when flexible deliveries are not allowed at all, but ADI is still available.

Additionally, we are interested in the question of whether the cost can be reduced even more if all orders are dispatched regardless of the remaining transportation capacity. Thus, the primary transportation option may already be exhausted, but additional orders will also be shipped to the facility using the alternative option. Our analysis can be easily adapted to obtain formulas for the computation of the expected total cost because the remaining units will only occur when the warehouse is running out of stock.

In Figure 4.6, we depict the average total cost for all three policies as a function of L_d. It can be seen that a shipment policy without flexible deliveries performs worst. For such a policy, we observe an almost linear cost decrease with an increasing demand lead time of approximately 0.7 % due to a reduction in the effective supply lead time, resulting in less safety stock at the warehouse.

Significant cost savings can be obtained with the introduction of flexible deliveries. However, shipping all orders one shipment day ahead and neglecting the available capacity can further be improved by our shipment policy, which takes the reserved capacity into account when deciding on the shipment quantity. Then, an increase in the demand lead time from 0 to 2 results in a decrease of the average total cost by more than 10 %, indicating that a reduction of approximately 9.5 % is caused by adapting the time-based

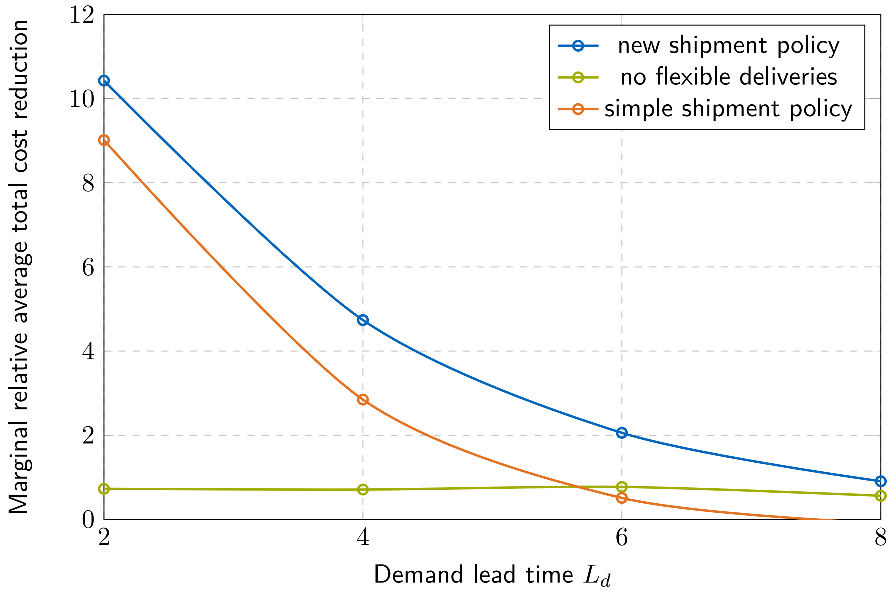

Figure 4.7: Marginal relative total cost reduction across all instances

shipment consolidation policy and allowing for flexible deliveries. We can conclude that the more significant part of the cost reduction is induced by the flexible delivery option and not by reducing the safety stock.

With increasing L_d, this effect is reduced until we reach a point at which additional ADI will only decrease the stock because the flexible delivery option is used to its extent. This can be seen in Figure 4.7 , where we illustrate the marginal relative cost reduction of all three policies.

We can also observe that the cost difference between the two flexible delivery concepts is increasing with increasing demand lead time. The marginal relative total cost reduction for the simple policy is even lower than for the shipment policy without flexible deliveries because, for large demand lead times, the alternative transportation option has to be used too often, which prevents further cost reductions.

Optimal Length of the Shipment Interval

Moreover, to enable early deliveries, even more safety stock is kept at the warehouse compared to a situation without flexible deliveries. This situation is illustrated in Table 4.8, where the optimal policy parameters R^* and T^* are presented for different values of the reserved transportation capacity, demand lead time, and order rate. The remaining

Table 4.8: Optimal reorder level and optimal shipment interval

		(R^*, T^*) shipment policy					(R^*, T^*) without flexible deliveries				
Cap	λ	$L_d = 0$	$L_d = 2$	$L_d = 4$	$L_d = 6$	$L_d = 8$	$L_d = 0$	$L_d = 2$	$L_d = 4$	$L_d = 6$	$L_d = 8$
5	1	(8,5)	(7,5)	(5,5)	(4,5)	(2,5)	(8,5)	(6,5)	(4,5)	(2,5)	(0,5)
	2	(20,3)	(16,3)	(12,3)	(7,3)	(3,3)	(20,3)	(15,3)	(11,3)	(6,3)	(2,3)
	4	(41,2)	(37,1)	(29,1)	(20,1)	(11,1)	(41,2)	(32,2)	(23,2)	(14,2)	(8,1)
10	1	(6,9)	(6,8)	(5,8)	(3,9)	(2,9)	(6,9)	(4,8)	(2,8)	(0,8)	(-2,8)
	2	(16,5)	(14,5)	(11,5)	(8,5)	(4,5)	(16,5)	(12,5)	(8,5)	(4,5)	(0,5)
	4	(37,3)	(30,3)	(22,3)	(14,3)	(6,3)	(37,3)	(29,3)	(20,3)	(12,3)	(4,3)
20	1	(6,15)	(5,15)	(5,15)	(4,15)	(3,15)	(6,15)	(4,15)	(2,15)	(0,15)	(-2,15)
	2	(17,9)	(15,9)	(14,9)	(11,9)	(7,9)	(17,9)	(12,9)	(8,9)	(4,9)	(0,9)
	4	(37,5)	(34,5)	(26,5)	(17,5)	(12,4)	(37,5)	(29,5)	(20,5)	(12,5)	(4,5)

parameters are fixed at $h = 1, w = 2, e = 2, c_1 = 20, c_2 = 2\frac{\omega}{Cap}$, and $L_s = 10$. While the demand lead time L_d has a significant impact on the numerical value of the reorder level, the influence on the optimal shipment interval is much less. This condition holds for both situations, with and without flexible deliveries. The optimal shipment interval is influenced by the shipment costs, as well as by the relationship between the average demand and the capacity. An increasing order rate raises the reorder level and reduces the shipment interval to utilize the reserved transportation capacity and avoid expensive additional shipments.

Optimal Transportation Capacity

We are further interested in the influence of flexible deliveries on the optimal capacity of the primary transportation option for a given shipment interval. In Table 4.9, it can be seen that flexible deliveries lead to equal or larger reorder levels for a fixed and given shipment interval. This means that more safety stock is required to enable flexible deliveries. Additionally, more reserved transportation capacity Cap^* is needed in situations with large values of T to exploit the benefit of flexible deliveries fully. However, the influence of flexible deliveries on the optimal capacity to be reserved is negligible.

4.4.3 Impact of the Policy Assumption

When we defined the flexible shipment consolidation policy, we only allowed orders to be shipped one shipment day ahead. They have to be shipped if enough stock and the remaining reserved transportation capacity is available. This assumption could be limiting

Table 4.9: Optimal reorder level and optimal capacity of the primary transportation option

		(R^*, Cap^*) with flexible deliveries					(R^*, Cap^*) without flexible deliveries				
T	λ	$L_d = 0$	$L_d = 2$	$L_d = 4$	$L_d = 6$	$L_d = 8$	$L_d = 0$	$L_d = 2$	$L_d = 4$	$L_d = 6$	$L_d = 8$
3	1	(9,3)	(8,3)	(7,3)	(4,3)	(2,3)	(9,3)	(7,3)	(5,3)	(2,3)	(0,3)
	2	(20,6)	(19,6)	(15,6)	(11,6)	(7,6)	(20,6)	(15,6)	(11,6)	(6,6)	(2,6)
	4	(38,11)	(38,12)	(25,12)	(20,12)	(10,12)	(38,11)	(30,11)	(22,12)	(14,12)	(6,12)
5	1	(8,5)	(7,5)	(6,5)	(4,5)	(2,5)	(8,5)	(6,5)	(4,5)	(2,5)	(0,5)
	2	(16,10)	(17,10)	(15,10)	(11,10)	(6,10)	(16,10)	(12,10)	(8,10)	(4,10)	(0,10)
	4	(37,20)	(37,21)	(23,20)	(20,20)	(21,21)	(37,20)	(29,20)	(20,20)	(12,20)	(4,20)
10	1	(6,10)	(6,10)	(6,11)	(5,11)	(4,11)	(6,10)	(4,10)	(2,10)	(0,10)	(-2,10)
	2	(16,20)	(17,21)	(16,22)	(16,22)	(14,22)	(16,20)	(12,20)	(8,20)	(4,20)	(0,20)
	4	(37,40)	(36,42)	(36,43)	(37,43)	(40,42)	(37,40)	(29,40)	(20,40)	(12,40)	(4,40)

in situations where L_d is much larger than T because orders could potentially be shipped several shipment days in advance. We want to investigate the impact of this assumption on the optimal decision and expected total cost. Therefore, we compare our policy with a policy where early shipments are always allowed if stock and capacity are available. That means orders are shipped as early as possible. For a fair comparison, we have to determine the optimal reorder level and the optimal shipment interval, which is done by a simulation-based approach.

We concentrate on a large demand lead time ($L_d = 8$) and a large demand rate ($\lambda = 4$) because, in these situations, it is more likely that the optimal shipment interval is smaller than L_d. L_s is fixed to 10 to reflect make-to-stock situations, and different reserved capacities are considered.

We expect a larger impact if late-delivery and early-delivery costs are high and therefore select the following cost parameters to test our conjecture: $l \in \{2, 10, 100\}$ and $e \in \{2, 10, 100\}$. The remaining parameters correspond to the base case.

The Table 4.10 columns two and three show the optimal decisions with and without the shipment assumption, while the last column presents the relative total cost deviation according to

$$\delta TC_{asm} = \frac{TC_{sim}(R^*_{asm}, T^*_{asm}) - TC_{sim}(R^*, T^*)}{TC_{sim}(R^*, T^*)} 100, \tag{4.28}$$

where $TC_{sim}(R^*, T^*)$ represents the minimal total cost without the shipment assumption and $TC_{sim}(R^*_{asm}, T^*_{asm})$ the minimal total cost with shipment assumption both determined

Table 4.10: Influence of shipment assumption on optimal reorder level and optimal shipment interval

		(R^*, T^*) with shipment assumption			(R^*, T^*) without shipment assumption			Relative total cost deviation		
Cap	l	$e = 2$	$e = 10$	$e = 100$	$e = 2$	$e = 10$	$e = 100$	$e = 2$	$e = 10$	$e = 100$
5	2	(10,2)	(6,3)	(4,3)	(9,2)	(6,3)	(4,3)	-0.5647	-0.0191	0.0107
	10	(15,2)	(10,2)	(8,3)	(14,2)	(9,2)	(8,3)	-1.8850	-2.8433	-0.4990
	100	(18,2)	(18,2)	(11,1)	(20,2)	(16,2)	(9,2)	4.6285	-15.7399	-14.5374
10	2	(8,2)	(5,3)	(2,3)	(7,2)	(5,3)	(2,3)	-0.7968	-0.0074	-0.1349
	10	(13,2)	(8,2)	(6,3)	(12,2)	(7,2)	(6,3)	-2.1444	-4.5346	-0.6919
	100	(18,2)	(18,2)	(10,1)	(19,2)	(14,2)	(7,2)	4.6304	-18.5438	-21.7269
20	2	(8,2)	(4,3)	(2,3)	(5,3)	(4,3)	(2,3)	-1.7276	-0.0583	-0.2639
	10	(11,2)	(8,2)	(5,3)	(10,2)	(7,2)	(4,3)	-2.1704	-6.6513	-0.5613
	100	(17,2)	(16,2)	(9,1)	(18,2)	(13,2)	(5,2)	4.3675	-19.8509	-23.2847

by simulation.

Our numerical experiments reveal that the shipment assumption impacts the optimal reorder level. If early shipments are allowed only one shipment day ahead, then more safety stock is needed at the warehouse to enable flexible deliveries in many situations. On the other hand, if early shipments are not restricted, then the early-delivery costs are controlled by a reduction of the optimal reorder level, making early deliveries less likely as stock-outs occur more frequently.

A contrary effect can be observed if the late-delivery costs parameter is very high ($l = 100$) and the early-delivery costs parameter is very low ($e = 2$). In such a situation, early deliveries are preferred to avoid waiting times, which results in larger safety stocks.

Our numerical results also show that the optimal shipment interval is relatively robust against the shipment assumption. Thus, total cost differences are induced by the different safety stock quantities at the warehouse. Although the reorder level can control the number of early deliveries, the opportunities are limited. This explains the lower minimal total cost when early deliveries are permitted only one shipment day ahead. Only in situations where early deliveries are much cheaper than late-delivery costs is it more beneficial to allow shipments as early as possible. However, it is unlikely that such large differences occur in reality since early-delivery and late-delivery costs are related to holding costs and, therefore, pretty much the same.

Based on our numerical study, we can conclude that, in line with the existing literature, ADI can lead to large cost reductions in inventory management. However, to fully exploit the benefit of ADI, the shipment policy should also be adapted, and flexible deliveries should be integrated into shipment consolidation programs. Doing so could entail larger inventories, but the savings due to a more efficient transportation policy far exceed the cost increases due to larger safety stocks.

4.5 Summary and Outlook

In this chapter, we have investigated a single-echelon inventory model with ADI and a flexible time-based shipment consolidation program with a reserved transportation capacity. We derive approximate mathematical expressions to compute shipment and inventory costs at the warehouse and thus are able to determine the warehouse reorder level and the shipment consolidation cycle length for the given situation. We have shown in a simulation study that our approximations have excellent performance and can be used to determine near-optimal policy parameters because the optimal decisions are found in more than 90 % of our instances, and the average total cost deviation is 0.1988 %.

The main finding is that companies can benefit greatly from ADI in the context of inventory management. However, they will miss opportunities if they only focus on the reduction of safety stocks and on logistic processes separately. Cost reductions can be further enhanced when logistic processes are connected, such as the shipment consolidation policy with flexible deliveries. In the investigated setting, the largest part of the cost reduction is induced by the flexible delivery option. Thus, the full potential of ADI can only be exploited if whole logistic processes are adapted.

Although not extensively explored in this chapter, we believe that our integrated logistic approach not only reduces costs but also has environmental benefits by increasing the utilization of the reserved transportation capacity. Further research can elaborate on these environmental aspects in more detail. It would be very interesting to understand how the optimal policy parameters and the optimal reserved transportation capacity behave if, besides the minimization of cost, the minimization of carbon emissions is also an aim.

We have assumed perfect ADI and identical demand lead times in our model. A logical next step is to replace the limiting assumptions and allow imperfect ADI as well as non-identical demand lead times. This increases the complexity of the model and requires a completely new analysis; therefore, it was not possible to investigate these aspects within the scope of this work.

Another direction for future research is an extension of the inventory system. Instead of only studying one warehouse, a divergent inventory system can be the object of future research. Another extension can be a more general demand model, such as a compound Poisson process.

However, we are convinced that these extensions will not change the main finding of this chapter, that ADI should not only be used to reduce stock levels but also to adapt related logistic processes. We expect all extensions to come with more complexity and require a heuristic solution approach.

4.6 Appendix

4.6.1 Derivation of the Probability Mass Function of the Shipment Quantity

We have to consider eleven cases to evaluate the shipment costs for various parameters. As stated in Section 4.2.1, the shipment quantity, $M(t)$, is always the sum of demands/orders during t_{n-1} and t_n plus the remaining units at t_{n-1}, $K(t_{n-1})$, minus the remaining units at t_n, $K(t_n)$. However, the formulas for the remaining units differ from case to case. In the following, we present the graphics and explanations considering the cases not discussed in Section 4.2.1.

Case 2 $L_d \leq T$ **and** $T < L_s \leq T + L_d$**:**

The main difference to Case 1 is that it cannot be assured that backorders are satisfied on the subsequent shipment day. Based on Figure 4.8, the formula for $K(t_n)$ can be expressed by

$$\begin{aligned} K(t_n) = \max\Bigg(\Bigg(&\underset{(R,Q)}{\operatorname{mod}} \Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-2}) - D(t_{n-2}, t_{n-1} - L_d) \\ &- D(t_{n-1} - L_d, t_n - L_s)\Big) - D(t_n - L_s, t_{n-1}) - D(t_{n-1}, t_n - L_d) \\ &- D(t_n - L_d, t_n)\Bigg)^-, \min\Big(D(t_n - L_d, t_n), \Big(Cap - D(t_{n-1}, t_n - L_d) \\ &- D(t_n - L_d, t_n) - K(t_{n-1})\Big)^-\Big)\Bigg), \end{aligned} \tag{4.29}$$

whereas the remaining units at t_{n-1} can be computed by

$$\begin{aligned} K(t_{n-1}) = \max\Bigg(&\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-2}) - D(t_{n-2}, t_{n-1} - L_d) \\ &- D(t_{n-1} - L_d, t_n - L_s) - D(t_n - L_s, t_{n-1})\Big)^-, \\ &\min\Big(D(t_{n-1} - L_d, t_n - L_s) + D(t_n - L_s, t_{n-1}), \Big(Cap \\ &- D(t_{n-2}, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_s) - D(t_n - L_s, t_{n-1}) \\ &- K(t_{n-2})\Big)^-\Big)\Bigg). \end{aligned} \tag{4.30}$$

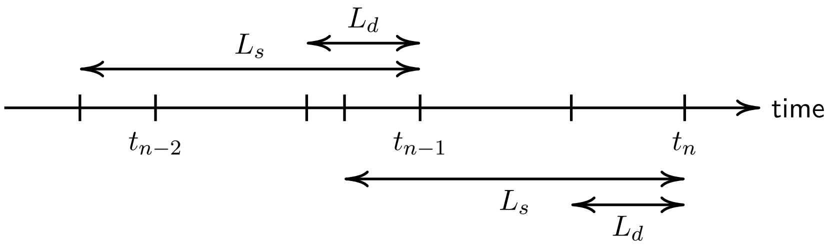

Figure 4.8: Shipment cycle when $L_d \leq T$ and $T < L_s \leq T + L_d$

Case 3 $L_d \leq T$ **and** $T + L_d < L_s \leq 2T$**:**

Compared to Case 2, $t_n - L_s$ is before $t_{n-1} - L_d$ which is illustrated in Figure 4.9. Based on the given sequence of time points, the remaining units at t_n can be obtained by

$$\begin{aligned} K(t_n) = \max \Bigg(\bigg(& \operatorname*{mod}_{R,Q} \Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-2}) - D(t_{n-2}, t_n - L_s) \Big) \\ & - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1}) - D(t_{n-1}, t_n - L_d) \\ & - D(t_n - L_d, t_n) \bigg)^-, \min \bigg(D(t_n - L_d, t_n), \Big(Cap - D(t_{n-1}, t_n - L_d) \\ & - D(t_n - L_d, t_n) - K(t_{n-1}) \Big)^- \bigg) \Bigg) \end{aligned} \tag{4.31}$$

and the remaining units at t_{n-1} by

$$\begin{aligned} K(t_{n-1}) = \max \bigg(& \Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-2}) - D(t_{n-2}, t_n - L_s) \\ & - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1}) \Big)^-, \\ & \min \Big(D(t_{n-1} - L_d, t_{n-1}), \big(Cap - D(t_{n-2}, t_n - L_s) \\ & - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1}) - K(t_{n-2}) \big)^- \Big) \bigg). \end{aligned} \tag{4.32}$$

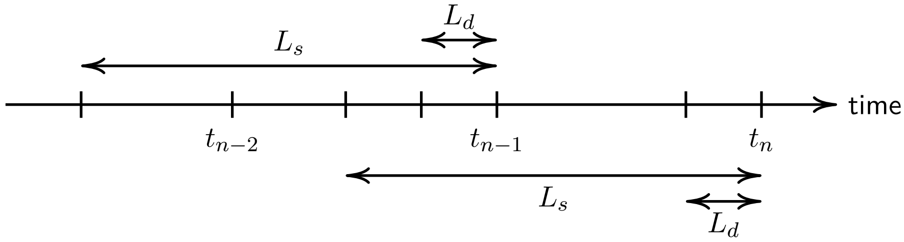

Figure 4.9: Shipment cycle when $L_d \leq T$ and $T + L_d < L_s \leq 2T$

Case 4 $L_d \leq T$ and $L_s > 2T$:

Case 4 is the last case where $L_d \leq T$ because as soon as $L_s > 2T$, the sequence of events is the same for each length of L_s. In Figure 4.10, an example is shown. The remaining units $K(t_n)$ can be computed by

$$\begin{aligned} K(t_n) = \max\Bigg(\Bigg(&\operatorname*{mod}_{(R,Q)} \Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s)\Big) - D(t_n - L_s, t_{n-2}) \\ &- D(t_{n-2}, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1}) - D(t_{n-1}, t_n - L_d) \\ &- D(t_n - L_d, t_n)\Bigg)^{-}, \min\Bigg(D(t_n - L_d, t_n), \Big(Cap - D(t_{n-1}, t_n - L_d) \\ &- D(t_n - L_d, t_n) - K(t_{n-1})\Big)^{-}\Bigg)\Bigg), \end{aligned} \tag{4.33}$$

whereas the remaining units at t_{n-1} can be obtained by

$$\begin{aligned} K(t_{n-1}) = \max\Bigg(\Big(&IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s) - D(t_n - L_s, t_{n-2}) \\ &- D(t_{n-2}, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1})\Big)^{-}, \min\Bigg(D(t_{n-1} - L_d, t_{n-1}), \\ &\Big(Cap - D(t_{n-2}, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-1}) - K(t_{n-2})\Big)^{-}\Bigg)\Bigg). \end{aligned} \tag{4.34}$$

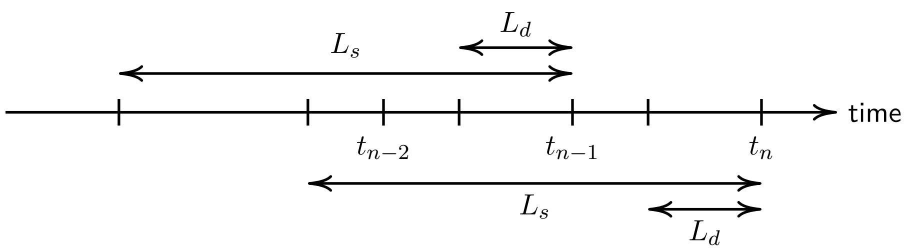

Figure 4.10: Shipment cycle when $L_d \leq T$ and $L_s > 2T$

Case 6 $T < L_d \leq 2T$ **and** $2T < L_s \leq T + L_d$**:**

Similar to Case 5, we now investigate cases where $T < L_d \leq 2T$, which means orders are known more than one shipment interval earlier. However, the shipment policy only allows units to be shipped to the production facility one shipment interval ahead. Therefore, all orders occurring during t_{n+1} and t_n cannot be shipped at t_n and are added to the remaining units directly. Considering the sequence of events given in Figure 4.11, we get

$$\begin{aligned} K(t_n) = D(t_{n+1} - L_d, t_n) + \max\Bigg(\bigg(& \underset{R,Q}{\operatorname{mod}} \Big(IP(t_{n-1} - L_s) \\ & - D(t_{n-1} - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_s)\Big) - D(t_n - L_s, t_{n-2}) \\ & - D(t_{n-2}, t_n - L_d) - D(t_n - L_d, t_{n-1}) - D(t_{n-1}, t_{n+1} - L_d)\bigg)^{-}, \\ & \min\Big(D(t_n - L_d, t_{n-1}) + D(t_{n-1}, t_{n+1} - L_d), \Big(Cap - D(t_{n-1}, t_{n+1} - L_d) \\ & - K(t_{n-1})\Big)^{-}\Big)\Bigg), \end{aligned} \tag{4.35}$$

and

$$\begin{aligned} K(t_{n-1}) = D(t_n - L_d, t_{n-1}) + \max\bigg(\Big(& IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\ & - D(t_{n-1} - L_d, t_n - L_s) - D(t_n - L_s, t_{n-2}) - D(t_{n-2}, t_n - L_d)\Big)^{-}, \\ & \min\Big(D(t_{n-1} - L_d, t_n - L_s) + D(t_n - L_s, t_{n-2}) + D(t_{n-2}, t_n - L_d), \\ & \Big(Cap - D(t_{n-2}, t_n - L_d) - K(t_{n-2})\Big)^{-}\Big)\bigg). \end{aligned} \tag{4.36}$$

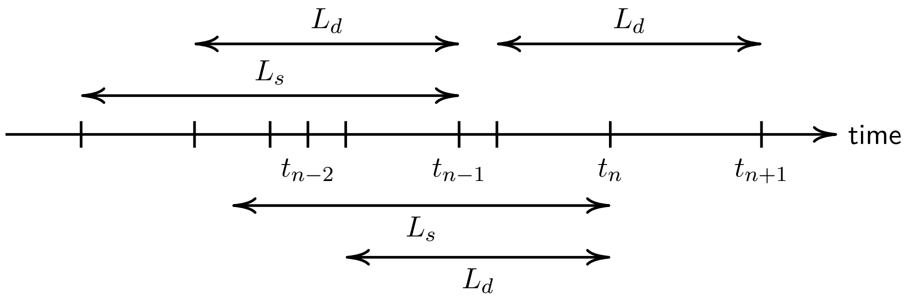

Figure 4.11: Shipment cycle when $T < L_d \leq 2T$ and $2T < L_s \leq T + L_d$

Case 7 $T < L_d \leq 2T$ **and** $L_s > T + L_d$**:**

Case 7 is close to Case 6; however, L_s has to be larger than $T + L_d$. In Figure 4.12, we can observe that $t_{n-1} - L_s$ and $t_n - L_s$ are the earliest time points, which is why an increase of L_s will not change the sequence of time points anymore. Therefore, we can obtain the remaining units at t_n by

$$K(t_n) = D(t_{n+1} - L_d, t_n) + \max\Bigg(\Big(\underset{R,Q}{\mathrm{mod}}\,\big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s)\big) - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-2}) - D(t_{n-2}, t_n - L_d) - D(t_n - L_d, t_{n-1}) - D(t_{n-1}, t_{n+1} - L_d)\Big)^-, \min\Big(D(t_n - L_d, t_{n-1}) + D(t_{n-1}, t_{n+1} - L_d), \big(Cap - D(t_{n-1}, t_{n+1} - L_d) - K(t_{n-1})\big)^-\Big)\Bigg). \tag{4.37}$$

and the remaining units at t_{n-1} by

$$K(t_{n-1}) = D(t_n - L_d, t_{n-1}) + \max\Bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s) - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_{n-2}) - D(t_{n-2}, t_n - L_d)\Big)^-, \min\Big(D(t_{n-1} - L_d, t_{n-2}) + D(t_{n-2}, t_n - L_d), \big(Cap - D(t_{n-2}, t_n - L_d) - K(t_{n-2})\big)^-\Big)\Bigg). \tag{4.38}$$

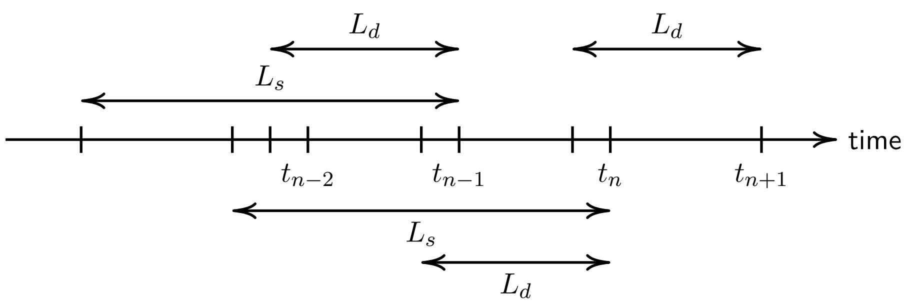

Figure 4.12: Shipment cycle when $T < L_d \leq 2T$ and $L_s > T + L_d$

Case 8 $2T < L_d \leq 3T$ **and** $2T < L_s \leq T + L_d$**:**

Now, we consider situations where $2T < L_d \leq 3T$. Case 8 is illustrated in Figure 4.13 which helps us to determine $K(t_n)$ and $K(t_{n-1})$. The remaining units at t_n can be computed by

$$K(t_n) = D(t_{n+1} - L_d, t_{n-1}) + D(t_{n-1}, t_n) + \max\Bigg(\bigg(\operatorname*{mod}_{R,Q}\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_s)\Big) - D(t_n - L_s, t_n - L_d) - D(t_n - L_d, t_{n-2}) - D(t_{n-2}, t_{n+1} - L_d)\bigg)^-, \min\Big(D(t_n - L_d, t_{n-2}) + D(t_{n-2}, t_{n+1} - L_d), \big(Cap + D(t_{n+1} - L_d, t_{n-1}) - K(t_{n-1})\big)^-\Big)\Bigg), \tag{4.39}$$

whereas the remaining units at t_n can be computed by

$$K(t_{n-1}) = D(t_n - L_d, t_{n-2}) + D(t_{n-2}, t_{n+1} - L_d) + D(t_{n+1} - L_d, t_{n-1}) + \max\Bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_s) - D(t_n - L_s, t_n - L_d)\Big)^-, \min\Big(D(t_{n-1} - L_d, t_n - L_s) + D(t_n - L_s, t_n - L_d), \big(Cap - K(t_{n-2})\big)^-\Big)\Bigg). \tag{4.40}$$

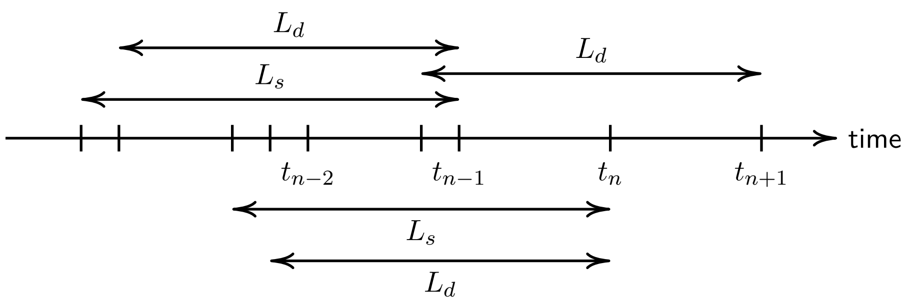

Figure 4.13: Shipment cycle when $2T < L_d \leq 3T$ and $2T < L_s \leq T + L_d$

Case 9 $2T < L_d \leq 3T$ **and** $L_s > T + L_d$**:**

Compared to Case 8, $t_n - L_s$ is before $t_{n-2} - L_d$ which is shown in Figure 4.14. Summarizing, the remaining units at t_n for Case 9 can be obtained by

$$\begin{aligned} K(t_n) = D(t_{n+1} - L_d, t_{n-1}) + D(t_{n-1}, t_n) + \max\Bigg(\Bigg(\underset{R,Q}{\text{mod}} \Big(IP(t_{n-1} - L_s) \\ - D(t_{n-1} - L_s, t_n - L_s)\Big) - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_d) \\ - D(t_n - L_d, t_{n-2}) - D(t_{n-2}, t_{n+1} - L_d)\Bigg)^{-}, \min\Big(D(t_n - L_d, t_{n-2}) \\ + D(t_{n-2}, t_{n+1} - L_d), \Big(Cap + D(t_{n+1} - L_d, t_{n-1}) - K(t_{n-1})\Big)^{-}\Big)\Bigg), \end{aligned} \tag{4.41}$$

and the remaining units at t_{n-1} by

$$\begin{aligned} K(t_{n-1}) = D(t_n - L_d, t_{n-2}) + D(t_{n-2}, t_{n+1} - L_d) + D(t_{n+1} - L_d, t_{n-1}) \\ + \max\Bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s) \\ - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_d)\Big)^{-}, \\ \min\Big(D(t_{n-1} - L_d, t_n - L_d), \Big(Cap - K(t_{n-2})\Big)^{-}\Big)\Bigg). \end{aligned} \tag{4.42}$$

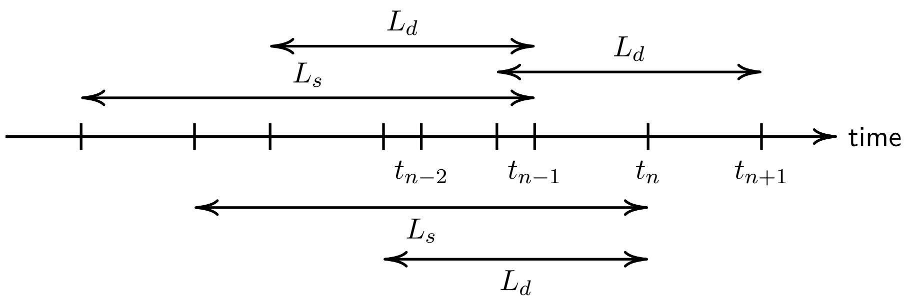

Figure 4.14: Shipment cycle when $2T < L_d \leq 3T$ and $L_s > T + L_d$

Case 10 $L_d > 3T$ **and** $3T < L_s \leq T + L_d$**:**

As soon as $L_s > 3T$, $t_{n+1} - L_d$ will always be before t_{n-2}, why this is the last range fore L_d. Again, we have to consider the length of L_s compared to $T + L_d$. In Figure 4.15, the considered case is illustrated, which leads to

$$\begin{aligned}
K(t_n) &= D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) + D(t_{n-1}, t_n) \\
&\quad + \max\Bigg(\bigg(\underset{R,Q}{\text{mod}}\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\
&\quad - D(t_{n-1} - L_d, t_n - L_s)\Big) - D(t_n - L_s, t_n - L_d) \\
&\quad - D(t_n - L_d, t_{n+1} - L_d)\bigg)^-, \min\Big(D(t_n - L_d, t_{n+1} - L_d), \\
&\quad \Big(Cap + D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) - K(t_{n-1})\Big)^-\Big)\Bigg),
\end{aligned} \tag{4.43}$$

and

$$\begin{aligned}
K(t_{n-1}) &= D(t_n - L_d, t_{n+1} - L_d) + D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) \\
&\quad + \max\bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_{n-1} - L_d) \\
&\quad - D(t_{n-1} - L_d, t_n - L_s) - D(t_n - L_s, t_n - L_d)\Big)^-, \\
&\quad \min\Big(D(t_{n-1} - L_d, t_n - L_s) + D(t_n - L_s, t_n - L_d), \Big(Cap \\
&\quad - K(t_{n-2})\Big)^-\Big)\bigg).
\end{aligned} \tag{4.44}$$

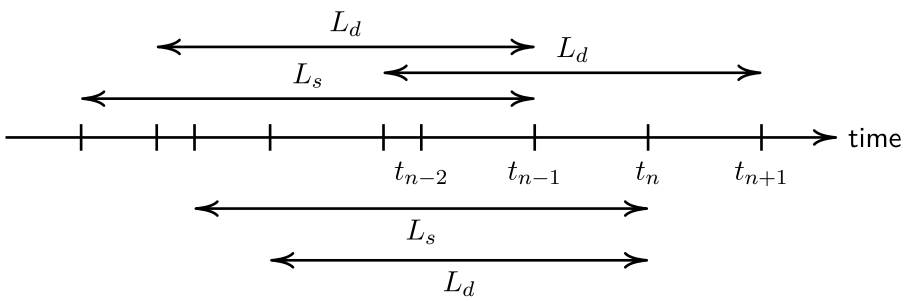

Figure 4.15: Shipment cycle when $L_d > 3T$ and $3T < L_s \leq T + L_d$

Case 11 $L_d > 3T$ **and** $L_s > T + L_d$**:**

Case 11 is shown in Figure 4.16 and represents the last case for the determination of the PMF of the shipment quantity. The remaining units at t_n can be computed by

$$\begin{aligned} K(t_n) = {} & D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) + D(t_{n-1}, t_n) \\ & + \max\Bigg(\bigg(\operatorname*{mod}_{R,Q}\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s)\Big) \\ & - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_d) \\ & - D(t_n - L_d, t_{n+1} - L_d)\bigg)^-, \min\Big(D(t_n - L_d, t_{n+1} - L_d), \Big(Cap \\ & + D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) - K(t_{n-1})\Big)^-\Big)\Bigg), \end{aligned} \tag{4.45}$$

and the remaining units at t_{n-1} by

$$\begin{aligned} K(t_{n-1}) = {} & D(t_n - L_d, t_{n+1} - L_d) + D(t_{n+1} - L_d, t_{n-2}) + D(t_{n-2}, t_{n-1}) \\ & + \max\bigg(\Big(IP(t_{n-1} - L_s) - D(t_{n-1} - L_s, t_n - L_s) \\ & - D(t_n - L_s, t_{n-1} - L_d) - D(t_{n-1} - L_d, t_n - L_d)\Big)^-, \\ & \min\Big(D(t_{n-1} - L_d, t_n - L_d), \Big(Cap - K(t_{n-2})\Big)^-\Big)\bigg). \end{aligned} \tag{4.46}$$

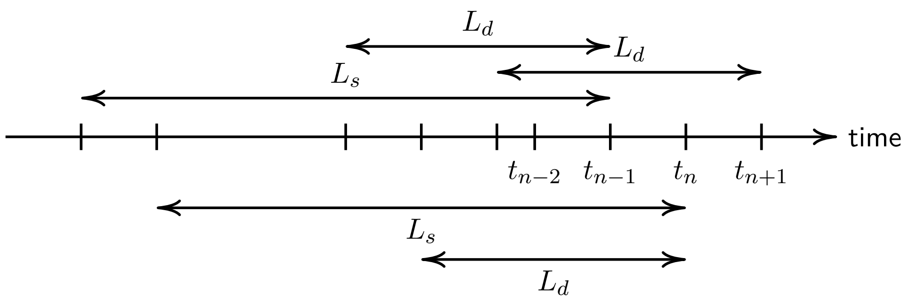

Figure 4.16: Shipment cycle when $L_d > 3T$ and $L_s > T + L_d$

4.6.2 Derivation of the Expected Inventory Costs when $S > 0$

In this section, we derive the cost expressions for the inventory costs for all possible situations. Therefore, we first determine the joint distribution function $f_i(x,y), i \in \{A,B,C,D,E,F,G\}$.

Situation A:

The derivation for $f_A(x,y)$ is already given in Section 4.2.2. Based on that, the inventory costs can be derived.

$$\begin{aligned}
E[C_A(\Omega(S),V)] &= \int_0^\infty \int_0^T C_A(x,y) f_A(x,y)\, dy\, dx \\
&= \int_{L_s}^\infty \int_0^{(T-L_d)^+} \Big(h(x - L_s + L_d + y) + ly\Big) g^S(x) u(y)\, dy\, dx \\
&= \int_{L_s}^\infty \frac{(T-L_d)^+}{T} \left(h\frac{S}{\lambda} g^{S+1}(x) + h(L_d - L_s) g^S(x) \right) \\
&\quad + (h+l)\frac{\Big((T-L_d)^+\Big)^2}{2T} g^S(x) dx \\
&= h\frac{(T-L_d)^+}{T}\left((L_d - L_s)\Big(1 - G^S(x)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big) \right) \\
&\quad + (h+l)\frac{\Big((T-L_d)^+\Big)^2}{2T}\Big(1 - G^S(L_s)\Big) \\
&= h\frac{(T-L_d)^+}{T}\left(\left(L_d - L_s + \frac{(T-L_d)^+}{2} \right)\Big(1 - G^S(L_s)\Big) \right. \\
&\quad \left. + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big) \right) + l\frac{\Big((T-L_d)^+\Big)^2}{2T}\Big(1 - G^S(L_s)\Big) \qquad (4.47)
\end{aligned}$$

The further derivation depends on the length of T and L_d, why we separate in $L_d \leq T$

and $L_d > T$. For $L_d \leq T$ we get

$$\begin{aligned}
E[C_A(\Omega(S), V)] &= h\frac{T - L_d}{T}\Bigg(\Bigg(L_d - L_s + \frac{T - L_d}{2}\Bigg)\big(1 - G^S(L_s)\big) \\
&\quad + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big) + l\frac{(T - L_d)^2}{2T}\big(1 - G^S(L_s)\big)\Bigg) \\
&= h\frac{T - L_d}{T}\Bigg(\Bigg(\frac{T + L_d}{2} - L_s\Bigg)\big(1 - G^S(L_s)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Bigg) \\
&\quad + l\frac{(T - L_d)^2}{2T}\big(1 - G^S(L_s)\big), \qquad (4.48)
\end{aligned}$$

whereas $L_d > T$ lead to

$$\begin{aligned}
E[C_A(\Omega(S), V)] &= h\frac{0}{T}\Bigg(\Bigg(L_d - L_s + \frac{0}{2}\Bigg)\big(1 - G^S(L_s)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Bigg) \\
&\quad + l\frac{0^2}{2T}\big(1 - G^S(L_s)\big) \\
&= 0. \qquad (4.49)
\end{aligned}$$

Summarizing, we get

$$E[C_A(\Omega(S), V)] = \begin{cases} h\frac{T-L_d}{T}\Big(\Big(\frac{T+L_d}{2} - L_s\Big)\big(1 - G^S(L_s)\big) \\ +\frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Big) \\ +l\frac{(T-L_d)^2}{2T}\big(1 - G^S(L_s)\big), & \text{if } L_d \leq T \\ \\ 0, & \text{otherwise.} \end{cases} \qquad (4.50)$$

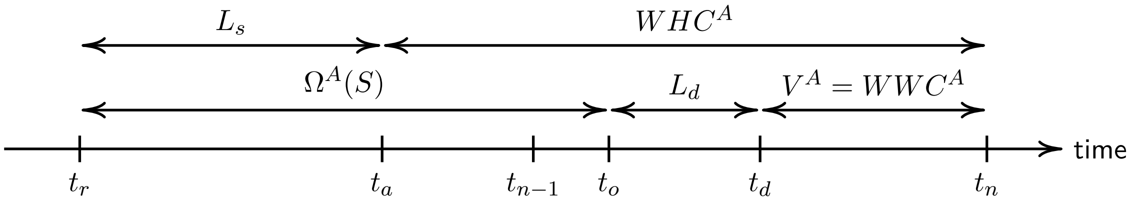

Figure 4.17: Timeline with the important time instances for Situation A when $S > 0$

Situation B:

To obtain the joint probability function $f_B(x,y)$, we determine the cumulative distribution function for Situation B by

$$F_B(x,y) = P\Big(\Omega(S) \leq x, V \leq y, t_a < t_o, t_o < t_{n-1} < t_d < t_n, t_s = t_{n-1}\big). \quad (4.51)$$

Thus,

$$\begin{aligned} F_B(x,y) &= P\Big(\Omega(S) \leq x, V \leq y, t_r + L_s < t_r + \Omega(S), \\ &\quad t_r + \Omega(S) < t_{n-1} < t_r + \Omega(S) + L_d < t_n\Big) \\ &= P\Big(\Omega(S) \leq x, V \leq y, L_s < \Omega(S), \\ &\quad \Omega(S) < t_{n-1} - t_r < \Omega(S) + L_d < t_n - t_r\Big). \end{aligned} \quad (4.52)$$

It holds that $T = t_r + \Omega(S) + L_d - t_{n-1} + V$, which leads to

$$\begin{aligned} F_B(x,y) &= P\Big(L_s < \Omega(S) \leq x, V \leq y, \\ &\quad \Omega(S) < \Omega(S) + L_d + V - T < \Omega(S) + L_d < t_n - t_{n-1} + t_{n-1} - t_r\Big) \\ &= P\Big(L_s < \Omega(S) \leq x, V \leq y, \\ &\quad \Omega(S) < \Omega(S) + L_d + V - T < \Omega(S) + L_d < \Omega(S) + L_d + V\Big) \\ &= P\Big(L_s < \Omega(S) \leq x, V \leq y, 0 < L_d + V - T < L_d < L_d + V\Big) \\ &= P\Big(L_s < \Omega(S) \leq x, (T - L_d)^+ < V \leq y\Big). \end{aligned} \quad (4.53)$$

For $x \geq L_s$ and $y \geq (T - L_d)^+$, we get

$$\begin{aligned} F_B(x,y) &= P\Big(L_s < \Omega(S) \leq x, V \leq y\Big) \\ &= \Big(G^S(x) - G^S(L_s)\Big)\Big(U(y) - U(T - L_d)\Big). \end{aligned} \quad (4.54)$$

Thus, the partial derivative with respect to both variables is given as

$$f_B(x,y) = \begin{cases} g^S(x)u(y) & L_s < x < \infty, (T - L_d)^+ < y \leq T \\ 0 & \text{otherwise.} \end{cases} \quad (4.55)$$

We get the expected costs

$$
\begin{aligned}
&E[C_B(\Omega(S),V)] \\
&= \int_0^\infty \int_0^T C_B(x,y) f_B(x,y)\, dy\, dx \\
&= \int_{L_s}^\infty \int_{(T-L_d)^+}^T \Big(hx + h(L_d - T - L_s) + eT + (h-e)y\Big) g^S(x) u(y)\, dy\, dx \\
&= \int_{L_s}^\infty \frac{T-(T-L_d)^+}{T} \Bigg(h\frac{S}{\lambda} g^{S+1}(x) + \Big(h(L_d - T - L_s) + eT\Big) g^S(x) \\
&\quad + (h-e)\frac{T^2 - \Big((T-L_d)^+\Big)^2}{2T} g^S(x)\Bigg)\, dx \\
&= \frac{T-(T-L_d)^+}{T} \Bigg(h\frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big) + \Big(h(L_d - T - L_s) + eT\Big) \cdot \\
&\quad \Big(1 - G^S(L_s)\Big) + (h-e)\frac{T^2 - \Big((T-L_d)^+\Big)^2}{2T}\Big(1 - G^S(L_s)\Big)\Bigg). \qquad (4.56)
\end{aligned}
$$

Assuming $L_d \leq T$, we get

$$
\begin{aligned}
&E[C_B(\Omega(S),V)] \\
&= \frac{L_d}{T}\Bigg(h\frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big) + \Big(h(L_d - T - L_s) + eT\Big)\Big(1 - G^S(L_s)\Big) \\
&\quad + (h-e)\frac{T^2 - (T-L_d)^2}{2T}\Big(1 - G^S(L_s)\Big)\Bigg) \\
&= h\frac{L_d}{T}\Bigg(\Bigg(L_d - T - L_s + \frac{2T - L_d}{2}\Bigg)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Bigg) \\
&\quad + e\frac{L_d}{T}\Bigg(T - \frac{2T - L_d}{2}\Bigg)\Big(1 - G^S(L_s)\Big) \\
&= h\frac{L_d}{T}\Bigg(\Bigg(\frac{L_d}{2} - L_s\Bigg)\Big(1 - G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\Bigg) \\
&\quad + e\frac{L_d^2}{2T}\Big(1 - G^S(L_s)\Big), \qquad (4.57)
\end{aligned}
$$

whereas for $L_d > T$, we get

$$\begin{aligned}
&E[C_B(\Omega(S), V)] \\
&= \frac{T-0}{T}\bigg(h\frac{S}{\lambda}\Big(1-G^{S+1}(L_s)\Big) + \Big(h(L_d - T - L_s) + eT\Big)\Big(1-G^S(L_s)\Big)\bigg) \\
&\quad + (h-e)\frac{T^2-0^2}{2T}\Big(1-G^S(L_s)\Big) \\
&= h\Bigg(\bigg(L_d - T - L_s + \frac{T}{2}\bigg)\Big(1-G^S(x)\Big) + \frac{S}{\lambda}\Big(1-G^{S+1}(L_s)\Big)\Bigg)+ \\
&\quad + e\bigg(T - \frac{T}{2}\bigg)\Big(1-G^S(L_s)\Big) \\
&= h\Bigg(\bigg(L_d - \frac{T}{2} - L_s\bigg)\Big(1-G^S(L_s)\Big) + \frac{S}{\lambda}\Big(1-G^{S+1}(L_s)\Big)\Bigg) \\
&\quad + e\frac{T}{2}\Big(1-G^S(L_s)\Big).
\end{aligned} \tag{4.58}$$

Finally we get

$$E[C_B(\Omega(S), V)] = \begin{cases} h\frac{L_d}{T}\Big(\Big(\frac{L_d}{2} - L_s\Big)\Big(1-G^S(L_s)\Big)\frac{S}{\lambda}\cdot \\ \Big(1-G^{S+1}(L_s)\Big)\Big) + e\frac{L_d^2}{2T}\Big(1-G^S(L_s)\Big) & \text{if } L_d \le T, \\ h\Big(\Big(L_d - \frac{T}{2} - L_s\Big)\Big(1-G^S(L_s)\Big) + \frac{S}{\lambda}\cdot \\ \Big(1-G^{S+1}(L_s)\Big)\Big) + e\frac{T}{2}\Big(1-G^S(L_s)\Big), & \text{otherwise.} \end{cases} \tag{4.59}$$

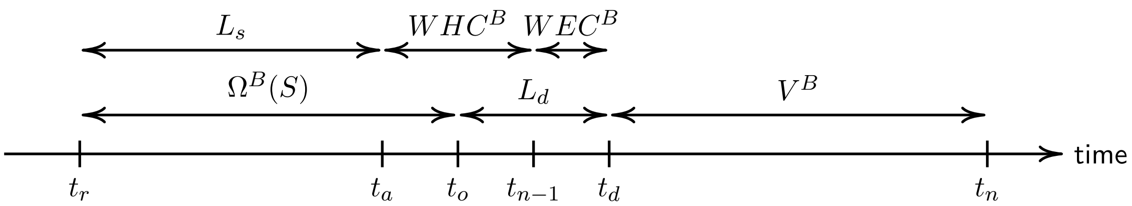

Figure 4.18: Timeline with the important time instances for Situation B when $S > 0$

Situation C:

Situation C is similar to Situation B except for the shipment point.

$$F_C(x,y) = P\Big(\Omega(S) \leq x, V \leq y, t_a < t_o, t_o < t_{n-1} < t_d < t_n, t_s = t_n\Big) \tag{4.60}$$

This is the situation with no early delivery.

$$\begin{aligned} F_C(x,y) = P\Big(&\Omega(S) \leq x, V \leq y, t_r + L_s < t_r + \Omega(S), \\ &t_r + \Omega(S) < t_{n-1} < t_r + \Omega(S) + L_d < t_n\Big) \end{aligned} \tag{4.61}$$

This yields

$$f_C(x,y) = \begin{cases} g^S(x)u(y) & L_s < x < \infty, (T-L_d)^+ \leq y \leq T, \\ 0 & \text{otherwise.} \end{cases} \tag{4.62}$$

The expected inventory costs for Situation C can be obtained by

$$\begin{aligned} E[C_C(\Omega(S),V)] &= \int_0^\infty \int_0^T C_C(x,y) f_C(x,y)\; dy\; dx \\ &= \int_{L_s}^\infty \int_{(T-L_d)^+}^T \Big(h(x - L_s + L_d + y) + ly\Big) g^S(x) u(y)\; dy\; dx \\ &= \int_{L_s}^\infty h\frac{T-(T-L_d)^+}{T}\left(\frac{S}{\lambda} g^{S+1}(x) + (L_d - L_s) g^S(x)\right) \\ &\quad + (h+l)\frac{T^2 - \Big((T-L_d)^+\Big)^2}{2T} g^S(x)\; dy\; dx \\ &= h\frac{T-(T-L_d)^+}{T}\left((L_d - L_s)\Big(1 - G^S(x)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\right) \\ &\quad + (h+l)\frac{T^2 - \Big((T-L_d)^+\Big)^2}{2T}\Big(1 - G^S(x)\Big). \end{aligned} \tag{4.63}$$

The length of T and L_d define the further cost derivation. For $L_d \leq T$ we get

$$\begin{aligned} E[C_C(\Omega(S),V)] = h&\frac{T-(T-L_d)}{T}\left((L_d - L_s)\Big(1 - G^S(x)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s)\Big)\right) \\ &+ (h+l)\frac{T^2 - (T-L_d)^2}{2T}\Big(1 - G^S(x)\Big) \end{aligned}$$

$$
\begin{aligned}
&= h\frac{L_d}{T}\Big((L_d - L_s)\big(1 - G^S(x)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Big) \\
&\quad + (h+l)\frac{2TL_d - L_d^2}{2T}\big(1 - G^S(x)\big) \\
&= h\frac{L_d}{T}\bigg(\Big(\frac{L_d}{2} - L_s + T\Big)\big(1 - G^S(L_s)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\bigg) \\
&\quad + l\frac{2TL_d - L_d^2}{2T}\big(1 - G^S(L_s)\big),
\end{aligned}
\tag{4.64}
$$

whereas $L_d > T$ occurs inventory costs of

$$
\begin{aligned}
E[C_C(\Omega(S), V)] &= h\frac{T-0}{T}\Big((L_d - L_s)\big(1 - G^S(x)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Big) \\
&\quad + (h+l)\frac{T^2 - 0^2}{2T}\big(1 - G^S(x)\big) \\
&= h\bigg(\Big(L_d - L_s + \frac{T}{2}\Big)\big(1 - G^S(x)\big) + \frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\bigg) \\
&\quad + l\frac{T}{2}\big(1 - G^S(x)\big).
\end{aligned}
\tag{4.65}
$$

Finally, the expected inventory costs for Situation C are defined as follows.

$$
E[C_C(\Omega(S), V)] =
\begin{cases}
h\frac{L_d}{T}\Big(\big(\frac{L_d}{2} - L_s + T\big)\big(1 - G^S(L_s)\big) \\
+\frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Big) \\
+l\frac{2TL_d - L_d^2}{2T}\big(1 - G^S(L_s)\big), & \text{if } L_d \leq T \\
h\Big(\big(L_d - L_s + \frac{T}{2}\big)\big(1 - G^S(L_s)\big) \\
+\frac{S}{\lambda}\big(1 - G^{S+1}(L_s)\big)\Big) \\
+l\frac{T}{2}\big(1 - G^S(L_s)\big) & \text{otherwise.}
\end{cases}
\tag{4.66}
$$

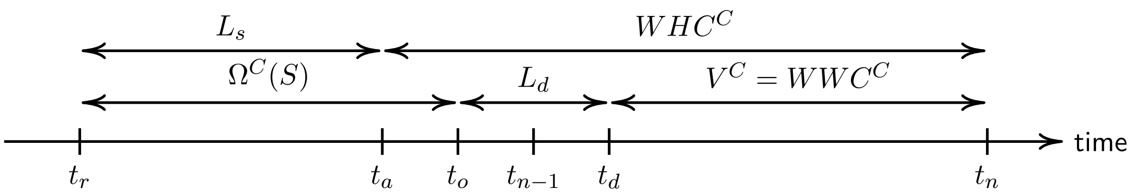

Figure 4.19: Timeline with the important time instances for Situation C when $S > 0$

Situation D:

Now, the considered uni is first ordered, then available, and then demanded.

$$\begin{aligned}
F_D(x,y) &= P\Big(\Omega(S) \le x, V \le y, t_o \le t_a,\\
&\qquad t_{n-1} < t_a < t_d < t_n, t_s = t_n\Big)\\
&= P\Big(\Omega(S) \le x, V \le y, t_r + \Omega(S) \le t_r + L_s,\\
&\qquad t_{n-1} < t_r + L_s < t_r + \Omega(S) + L_d < t_n\Big)\\
&= P\Big(\Omega(S) \le x, V \le y, \Omega(S) \le L_s,\\
&\qquad 0 < t_r + L_s - t_{n-1} < t_r + \Omega(S) + L_d - t_{n-1} < T\Big) \qquad (4.67)
\end{aligned}$$

Since $T = V + L_d + t_r + \Omega(S) - t_{n-1}$ we get

$$\begin{aligned}
F_D(x,y) &= P\Big(\Omega(S) \le x, \Omega(S) \le L_s, V \le y,\\
&\qquad 0 < T - V - L_d - \Omega(S) + L_s < T - V < T\Big)\\
&= P\Big(L_s - L_d \le \Omega(S) \le x, \Omega(S) < L_s, V \le y,\\
&\qquad V < T - L_d - \Omega(S) + L_s\Big) \qquad (4.68)
\end{aligned}$$

For $L_s - L_d \le x \le L_s$ and $y \le T - L_d + L_s - x$ we get

$$\begin{aligned}
F_D(x,y) &= P\Big(L_s - L_d \le \Omega(S) \le x, V \le y, V < T - L_d - \Omega(S) + L_s\Big)\\
&= \int_{L_s - L_d}^{x} P\Big(V \le y, V < T - L_d - t + L_s \mid \Omega(S) = t\Big) g^S(t)\, dt\\
&= \int_{L_s - L_d}^{x} P\Big(V \le \min(y, T - L_d + L_s - t)\Big) g^S(t)\, dt\\
&= \int_{L_s - L_d}^{x} \frac{1}{T} y g^S(t)\, dt\\
&= \frac{y}{T}\Big(G^S(x) - G^S(L_s - L_d)\Big). \qquad (4.69)
\end{aligned}$$

To obtain $f_D(x,y)$, we have to compute a partial derivative with respect to x and y. For $L_s - L_d \leq x < L_s$ and $y \leq T - L_d + L_s - x$ we get

$$\begin{aligned} f_D(x,y) &= \frac{\partial}{\partial x \partial y} F_D(x,y) \\ &= \frac{\partial}{\partial x \partial y}\Big(\frac{y}{T}\Big(G^S(x) - G^S(L_s - L_d)\Big)\Big) \\ &= \frac{\partial}{\partial y}\Big(\frac{y}{T} g^S(x)\Big) \\ &= \frac{1}{T} g^S(x) = g^S(x)u(y). \end{aligned} \tag{4.70}$$

Thus, we obtain

$$f_D(x,y) = \begin{cases} g^S(x)u(y) & L_s - L_d < x \leq L_s, 0 \leq y \leq (T - L_d + L_s - x)^+ \\ 0 & \text{otherwise.} \end{cases} \tag{4.71}$$

The inventory costs can be derived by

$$\begin{aligned} &E[C_D(\Omega(S), V)] \\ &= \int_0^\infty \int_0^T C_D(x,y) f_D(x,y)\; dy\; dx \\ &= \int_{L_s - L_d}^{L_s} \int_0^{(T-L_d+L_s-x)^+} \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\; dy\; dx \end{aligned} \tag{4.72}$$

Now, we have to separate the interval $L_s - L_d \leq x < L_s$ in $L_s - L_d \leq x \leq L_s - (L_d - T)^+$ and $L_s - (L_d - T)^+ < x < L_s$ to divide $0 \leq y \leq (T - L_d + L_s - x)^+$ into $0 \leq y \leq T - L_d + L_s - x$ and $0 \leq y \leq 0$. We get

$$\begin{aligned} &E[C_D(\Omega(S), V)] \\ &= \int_{L_s - L_d}^{L_s - (L_d - T)^+} \int_0^{T - L_d + L_s - x} \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\; dy\; dx \\ &\quad + \int_{L_s - (L_d - T)^+}^{L_s} \int_0^0 \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\; dy\; dx \\ &= \int_{L_s - L_d}^{L_s - (L_d - T)^+} h\frac{T - L_d + L_s - x}{T}\frac{S}{\lambda} g^{S+1}(x) + h(L_d - L_s)\frac{T - L_d + L_s - x}{T} g^S(x) \\ &\quad + (h+l)\frac{(T - L_d + L_s - x)^2}{2T} g^S(x)\; dx \end{aligned}$$

$$
\begin{aligned}
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\biggl(\frac{2(L_d-L_s)(T-L_d+L_s)}{2T} + \frac{(T-L_d+L_s)^2}{2T}\biggr) g^S(x) \\
&\quad + h\biggl(\frac{T-L_d+L_s}{T} - \frac{L_d-L_s}{T} - \frac{T-L_d+L_s}{T}\biggr)\frac{S}{\lambda} g^{S+1}(x) - h\frac{S(S+1)}{2T\lambda^2} g^{S+2}(x) \\
&\quad + l\frac{(T-L_d+L_s)^2}{2T} g^S(x) - l\frac{(T-L_d+L_s)S}{T\lambda} g^{S+1}(x) + l\frac{S(S+1)}{2T\lambda^2} g^{S+2}(x)\, dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\biggl(\frac{2TL_d - 2L_d^2 + 2L_sL_d - 2TL_s + 2L_sL_d - 2L_s^2 + T^2 + L_d^2 + L_s^2}{2T} \\
&\quad + \frac{-2TL_d + 2TL_s - 2L_sL_d}{2T}\biggr) g^S(x) + h\frac{(L_s-L_d)S}{T\lambda} g^{S+1}(x) - h\frac{S(S+1)}{2T\lambda^2} g^{S+2}(x) \\
&\quad + l\frac{(T-L_d+L_s)^2}{2T} g^S(x) - l\frac{(T-L_d+L_s)S}{T\lambda} g^{S+1}(x + l\frac{S(S+1)}{2T\lambda^2} g^{S+2}(x)\, dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\biggl(\frac{T^2-(L_s-L_d)^2}{2T} g^S(x) + \frac{(L_s-L_d)S}{T\lambda} g^{S+1}(x) \\
&\quad - \frac{S(S+1)}{2T\lambda^2} g^{S+2}(x)\biggr) + l\biggl(\frac{(T-L_d+L_s)^2}{2T} g^S(x) - \frac{(T-L_d+L_s)S}{T\lambda} g^{+1}S(x) \\
&\quad + \frac{S(S+1)}{2T\lambda^2} g^{S+2}(x)\biggr)\, dx \\
&= h\biggl(\frac{T^2-(L_s-L_d)^2}{2T}\Bigl(G^S(L_s-(L_d-T)^+) - G^S(L_s-L_d)\Bigr) \\
&\quad + \frac{(L_s-L_d)S}{T\lambda}\Bigl(G^{S+1}(L_s-(L_d-T)^+) - G^{S+1}(L_s-L_d)\Bigr) \\
&\quad - \frac{S(S+1)}{2T\lambda^2}\Bigl(G^{S+2}(L_s-(L_d-T)^+) - G^{S+2}(L_s-L_d)\Bigr)\biggr) \\
&\quad + l\biggl(\frac{(T-L_d+L_s)^2}{2T}\Bigl(G^S(L_s-(L_d-T)^+) - G^S(L_s-L_d)\Bigr) \\
&\quad - \frac{(T-L_d+L_s)S}{T\lambda}\Bigl(G^{S+1}(L_s-(L_d-T)^+) - G^{S+1}(L_s-L_d)\Bigr) \\
&\quad + \frac{S(S+1)}{2T\lambda^2}\Bigl(G^{S+2}(L_s-(L_d-T)^+) - G^{S+2}(L_s-L_d)\Bigr)\biggr). \qquad (4.73)
\end{aligned}
$$

Summarizing, we get

$$
\begin{aligned}
&E[C_D(\Omega(S),V)] \qquad (4.74)\\
&= \begin{cases}
h\Big(\frac{T^2-(L_s-L_d)^2}{2T}\Big(G^S(L_s)-G^S(L_s-L_d)\Big) \\
+\frac{(L_s-L_d)S}{T\lambda}\Big(G^{S+1}(L_s)-G^{S+1}(L_s-L_d)\Big) \\
-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s)-G^{S+2}(L_s-L_d)\Big)\Big) \\
+l\Big(\frac{(T-L_d+L_s)^2}{2T}\Big(G^S(L_s)-G^S(L_s-L_d)\Big) \\
-\frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s)-G^{S+1}(L_s-L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s)-G^{S+2}(L_s-L_d)\Big)\Big), & \text{if } L_d \leq T \\
\\
h\Big(\frac{T^2-(L_s-L_d)^2}{2T}\Big(G^S(L_s-L_d+T)-G^S(L_s-L_d)\Big) \\
+\frac{(L_s-L_d)S}{T\lambda}\Big(G^{S+1}(L_s-L_d+T)-G^{S+1}(L_s-L_d)\Big) \\
-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-L_d+T)-G^{S+2}(L_s-L_d)\Big)\Big) \\
+l\Big(\frac{(T-L_d+L_s)^2}{2T}\Big(G^S(L_s-L_d+T)-G^S(L_s-L_d)\Big) \\
-\frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s-L_d+T)-G^{S+1}(L_s-L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-L_d+T)-G^{S+2}(L_s-L_d)\Big)\Big), & \text{otherwise.}
\end{cases}
\end{aligned}
$$

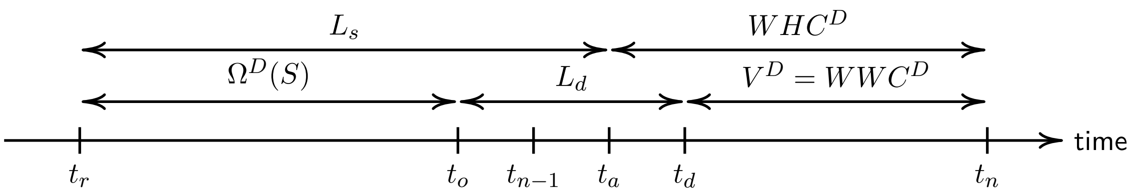

Figure 4.20: Timeline with the important time instances for Situation D when $S > 0$

Situation E:

Now, t_o and t_a occur in different shipment intervals, which is why flexible deliveries are possible if enough capacity is available.

$$\begin{aligned}F_E(x,y) &= P\Big(\Omega(S) \leq x, V \leq y, t_o \leq t_a, t_a < t_{n-1} < t_d < t_n, t_s = t_{n-1}\Big)\\ &= P\Big(\Omega(S) \leq x, V \leq y, t_r + \Omega(S) \leq t_r + L_s,\\ &\quad t_r + L_s < t_{n-1} < t_r + \Omega(S) + L_d < t_n\Big)\\ &= P\Big(\Omega(S) \leq x, V \leq y, \Omega(S) \leq L_s,\\ &\quad L_s < t_{n-1} - t_r < \Omega(S) + L_d < t_n - t_r\Big)\end{aligned} \tag{4.75}$$

Since $T = t_r + \Omega(S) + L_d - t_{n-1} + V$, we get $t_{n-1} - t_r = \Omega(S) + L_d - T + V$

$$\begin{aligned}F_E(x,y) &= P\Big(\Omega(S) \leq \min(x, L_s), V \leq y, L_s < \Omega(S) + L_d - T + V\\ &\quad < \Omega(S) + L_d < t_n - t_{n-1} + t_{n-1} - t_r, M(t_n) < Cap\Big)\\ &= P\Big(\Omega(S) \leq \min(x, L_s), V \leq y, L_s < \Omega(S) + L_d - T+\\ &\quad V < \Omega(S) + L_d < \Omega(S) + L_d + V, M(t_n) < Cap\Big)\\ &= P\Big(L_s - L_d < \Omega(S) \leq \min(x, L_s), T + L_s - \Omega(S) - L_d < V \leq y\Big).\end{aligned} \tag{4.76}$$

For $L_s - L_d \leq x \leq L_s$ we get

$$F_E(x,y) = \int_{L_s - L_d}^{x} P\Big(T + L_s - t - L_d < V \leq y\Big) g^S(t)\, dt. \tag{4.77}$$

For $T + L_s - L_d - x \geq y$, this is zero. Now $T + L_s - L_d - x < y$

$$
\begin{aligned}
F_E(x,y) &= \int_{T+L_s-L_d-y}^{x} \Big(P(V \leq y) - P(V \leq T + L_s - t - L_d)\Big) g^S(t)\ dt \\
&= \int_{T+L_s-L_d-y}^{x} \left(\frac{y}{T} - \frac{T + L_s - t - L_d}{T}\right) g^S(t)\ dt \\
&= \frac{y + L_d - T - L_s}{T}\Big(G^S(x) - G^S(T + L_s - L_d - y)\Big) \\
&\quad + \frac{1}{T}\int_{T+L_s-L_d-y}^{x} t g^S(t)\ dt \\
&= \frac{y + L_d - T - L_s}{T}\Big(G^S(x) - G^S(T + L_s - L_d - y)\Big) \\
&\quad + \frac{1}{T}\frac{S}{\lambda}\int_{T+L_s-L_d-y}^{x} g^{S+1}(t)\ dt \\
&= \frac{y + L_d - T - L_s}{T}\Big(G^S(x) - G^S(T + L_s - L_d - y)\Big) \\
&\quad + \frac{1}{T}\frac{S}{\lambda}\Big(G^{S+1}(x) - G^{S+1}(T + L_s - L_d - y)\Big).
\end{aligned} \tag{4.78}
$$

To obtain $f_E(x,y)$, we have to compute partial derivative with respect to x and y. For $L_s - L_d \leq x < L_s$ and $y > T - L_d + L_s - x$ we get

$$
\begin{aligned}
f_E(x,y) &= \frac{\partial}{\partial x \partial y} F_E(x,y) \\
&= \frac{\partial}{\partial x \partial y}\Bigg(\frac{y + L_d - T - L_s}{T}\Big(G^S(x) - G^S(T + L_s - L_d - y)\Big) \\
&\quad + \frac{1}{T}\frac{S}{\lambda}\Big(G^{S+1}(x) - G^{S+1}(T + L_s - L_d - y)\Big)\Bigg) \\
&= \frac{\partial}{\partial y}\left(\frac{y + L_d - T - L_s}{T} g^S(x) + \frac{1}{T}\frac{S}{\lambda} g^{S+1}(x)\right) \\
&= \frac{1}{T} g^S(x) = g^S(x) u(y).
\end{aligned} \tag{4.79}
$$

Thus, we obtain

$$
f_E(x,y) = \begin{cases} g^S(x)u(y) & L_s - L_d < x \leq L_s, T - L_d + L_s - x < y \leq T \\ 0 & \text{otherwise.} \end{cases} \tag{4.80}
$$

Now, we can derive the inventory costs as follows.

$$\begin{aligned}
&E[C_E(\Omega(S),V)]\\
&=\int_0^\infty\int_0^T C_E(x,y)f_E(x,y)\;dy\;dx\\
&=\int_{L_s-L_d}^{L_s}\int_{(T-L_d+L_s-x)^+}^{T}\Big(h(x-T-L_s+L_d+y)+e(T-y)\Big)g^S(x)u(y)\;dy\;dx\\
&=\int_{L_s-L_d}^{L_s}\int_{(T-L_d+L_s-x)^+}^{T}\Big(hx+h(L_d-T-L_s)+eT+(h-e)y\Big)g^S(x)u(y)\;dy\;dx
\end{aligned}\tag{4.81}$$

Again, we have to separate the interval $L_s-L_d\le x<L_s$ in $L_s-L_d\le x\le L_s-(L_d-T)^+$ and $L_s-(L_d-T)^+<x<L_s$ to divide $(T-L_d+L_s-x)^+<y\le T$ into $T-L_d+L_s-x<y\le T$ and $0\le y\le T$.

$$\begin{aligned}
&E[C_E(\Omega(S),V)]\\
&=\int_{L_s-L_d}^{L_s-(L_d-T)^+}\int_{T-L_d+L_s-x}^{T}\Big(hx+h(L_d-T-L_s)+eT+(h-e)y\Big)g^S(x)u(y)\;dy\;dx\\
&\quad+\int_{L_s-(L_d-T)^+}^{L_s}\int_0^T\Big(hx+h(L_d-T-L_s)+eT+(h-e)y\Big)g^S(x)u(y)\;dy\;dx\\
&=\int_{L_s-L_d}^{L_s-(L_d-T)^+} h\frac{L_d-L_s+x}{T}\frac{S}{\lambda}g^{S+1}(x)+h(L_d-T-L_s)\frac{L_d-L_s+x}{T}g^S(x)\\
&\quad+eT\frac{L_d-L_s+x}{T}g^S(x)+(h-e)\frac{T^2-(T-L_d+L_s-x)^2}{2T}g^S(x)\;dx\\
&\quad+\int_{L_s-(L_d-T)^+}^{L_s} h\frac{S}{\lambda}g^{S+1}(x)+h(L_d-T-L_s)g^S(x)+eTg^S(x)+(h-e)\frac{T}{2}g^S(x)\;dx\\
&=\int_{L_s-L_d}^{L_s-(L_d-T)^+} h\frac{L_d-L_s}{T}\frac{S}{\lambda}g^{S+1}(x)+h\frac{S(S+1)}{T\lambda^2}g^{S+2}(x)\\
&\quad+h\frac{(L_d-T-L_s)(L_d-L_s)}{T}g^S(x)+h\frac{(L_d-T-L_s)}{T}\frac{S}{\lambda}g^{S+1}(x)\\
&\quad+e\frac{T(L_d-L_s)}{T}g^S(x)+e\frac{TS}{T\lambda}gS+1(x)+(h-e)\frac{T^2-(T-L_d+L_s)^2}{2T}g^S(x)\;dx\\
&\quad+(h-e)\frac{2(T-L_d+L_s)}{2T}\frac{S}{\lambda}g^{S+1}(x)-(h-e)\frac{S(S+1)}{T\lambda^2}g^{S+2}(x)+\\
&\quad\int_{L_s-(L_d-T)^+}^{L_s} h\left(L_d-L_s-\frac{T}{2}\right)g^S(x)+h\frac{S}{\lambda}g^{S+1}(x)+e\frac{T}{2}g^S(x)\;dx\\
&=\int_{L_s-L_d}^{L_s-(L_d-T)^+} h\left(\frac{2(L_d-T-L_s)(L_d-L_s)}{2T}+\frac{T^2-(T-L_d+L_s)^2}{2T}\right)g^S(x)\\
&\quad+h\left(\frac{L_d-L_s}{T}+\frac{L_d-T-L_s}{T}+\frac{T-L_d+L_s}{T}\right)\frac{S}{\lambda}g^{S+1}(x)+h\frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)
\end{aligned}$$

$$
\begin{aligned}
&+ e\bigg(\frac{2T(L_d-L_s)}{2T} - \frac{T^2-(T-L_d+L_s)^2}{2T}\bigg)g^S(x) + e\frac{S(L_d-L_s)}{T\lambda}g^{S+1}(x) \\
&+ e\frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\,dx + \int_{L_s-(L_d-T)^+}^{L_s} h(L_d-L_s-\frac{T}{2})g^S(x) + h\frac{S}{\lambda}g^{S+1}(x) \\
&+ e\frac{T}{2}g^S(x)\,dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\Bigg(\bigg(\frac{2L_d^2-2TL_d-2L_sL_d-2L_sL_d+2TL_s+2L_s^2+T^2-T^2-L_d^2}{2T} \\
&+ \frac{-L_s^2+2TL_d-2TL_s+2L_sL_d}{2T}\bigg)g^S(x) + \frac{(L_s-L_d)S}{T\lambda}g^{S+1}(x) \\
&- \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\Bigg) + e\Bigg(\bigg(\frac{2TL_d-2TL_s-T^2+T^2+L_d^2+L_s^2-2TL_d+2TL_s}{2T} \\
&+ \frac{-2L_sL_d}{2T}\bigg)g^S(x) + \frac{(L_d-L_s)S}{T\lambda}g^{S+1}(x) + \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\Bigg)\,dx \\
&+ \int_{L_s-(L_d-T)^+}^{L_s} h\Bigg(\bigg(L_d-L_s-\frac{T}{2}\bigg)g^S(x) + \frac{S}{\lambda}g^{S+1}(x)\Bigg) + e\frac{T}{2}g^S(x)\,dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\bigg(\frac{(L_d-L_s)^2}{2T}g^S(x) + \frac{(L_d-L_s)S}{T\lambda}g^{S+1}(x) + \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\bigg) \\
&+ e\bigg(\frac{(L_d-L_s)^2}{2T}g^S(x) + \frac{(L_d-L_s)S}{T\lambda}g^{S+1}(x) + \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\bigg)\,dx \\
&+ \int_{L_s-(L_d-T)^+}^{L_s} h\Bigg(\bigg(L_d-L_s-\frac{T}{2}\bigg)g^S(x) + \frac{S}{\lambda}g^{S+1}(x)\Bigg) + e\frac{T}{2}g^S(x)\,dx \\
&= h\bigg(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s-(L_d-T)^+) - G^S(L_s-L_d)\Big) \\
&+ \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s-(L_d-T)^+) - G^{S+1}(L_s-L_d)\Big) \\
&+ \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-(L_d-T)^+) - G^{S+2}(L_s-L_d)\Big)\bigg) \\
&+ e\bigg(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s-(L_d-T)^+) - G^S(L_s-L_d)\Big) \\
&+ \frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s-(L_d-T)^+) - G^{S+1}(L_s-L_d)\Big) \\
&+ \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-(L_d-T)^+) - G^{S+2}(L_s-L_d)\Big)\bigg) \\
&+ h\Bigg(\bigg(L_d-L_s-\frac{T}{2}\bigg)\Big(G^S(L_s) - G^S(L_s-(L_d-T)^+)\Big)
\end{aligned}
$$

$$
+ \frac{S}{\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - (L_d - T)^+)\Big)\Bigg)
$$

$$
+ e\frac{T}{2}\Big(G^{S}(L_s) - G^{S}(L_s - (L_d - T)^+)\Big) \tag{4.82}
$$

Summarizing we get

$$
E[C_E(\Omega(S), V)] \tag{4.83}
$$

$$
= \begin{cases}
h\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) \\
+\frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big) \\
+e\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s) - G^S(L_s - L_d)\Big) \\
+\frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s) - G^{S+2}(L_s - L_d)\Big)\Big) & \text{if } L_d \leq T, \\
\\
h\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) \\
+\frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) \\
+(L_d - L_s - \frac{T}{2})\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big) \\
+\frac{S}{\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - L_d + T)\Big)\Big) \\
+e\Big(\frac{(L_d-L_s)^2}{2T}\Big(G^S(L_s - L_d + T) - G^S(L_s - L_d)\Big) \\
+\frac{(L_d-L_s)S}{T\lambda}\Big(G^{S+1}(L_s - L_d + T) - G^{S+1}(L_s - L_d)\Big) \\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - L_d + T) - G^{S+2}(L_s - L_d)\Big) \\
+\frac{T}{2}\Big(G^S(L_s) - G^S(L_s - L_d + T)\Big)\Big) & \text{otherwise.}
\end{cases}
$$

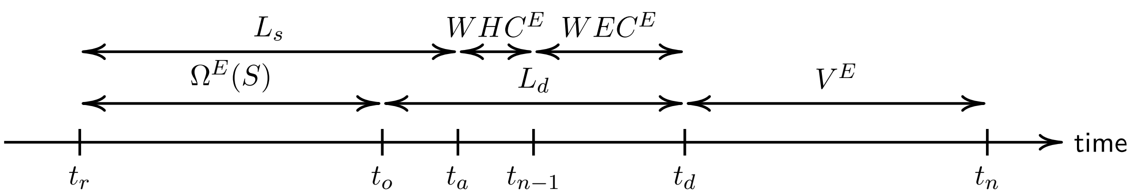

Figure 4.21: Timeline with the important time instances for Situation E when $S > 0$

Situation F:

The inventory costs changes if no reserved transportation capacity is available for the considered unit.

$$F_F(x,y) = P\Big(\Omega(S) \le x, V \le y, t_o \le t_a, t_a < t_{n-1} < t_d < t_n, t_s = t_n\Big) \tag{4.84}$$

We get the same probabilities since the shipment time is different only from that of Situation E. Thus, we obtain

$$f_F(x,y) = \begin{cases} g^S(x)u(y) & L_s - L_d < x \le L_s, T - L_d + L_s - x < y \le T \\ 0 & \text{otherwise.} \end{cases} \tag{4.85}$$

The inventory costs are computed as follows.

$$\begin{aligned} &E[C_F(\Omega(S), V)] \\ &= \int_0^\infty \int_0^T C_F(x,y) f_F(x,y)\, dy\, dx \\ &= \int_{L_s-L_d}^{L_s} \int_{(T-L_d+L_s-x)^+}^T \Big(h(x - L_s + L_d + y) + ly\Big) g^S(x)u(y)\, dy\, dx \\ &= \int_{L_s-L_d}^{L_s} \int_{(T-L_d+L_s-x)^+}^T \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\, dy\, dx \end{aligned} \tag{4.86}$$

The interval $L_s - L_d \le x < L_s$ has to be separated in $L_s - L_d \le x \le L_s - (L_d - T)^+$ and $L_s - (L_d - T)^+ < x < L_s$ to divide the integral $(T - L_d + L_s - x)^+ < y \le T$ into $T - L_d + L_s - x < y \le T$ and $0 \le y \le T$.

$$\begin{aligned} &E[C_F(\Omega(S), V)] \\ &= \int_{L_s-L_d}^{L_s-(L_d-T)^+} \int_{T-L_d+L_s-x}^T \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\, dy\, dx \\ &\quad + \int_{L_s-(L_d-T)^+}^{L_s} \int_0^T \Big(hx + h(L_d - L_s) + (h+l)y\Big) g^S(x)u(y)\, dy\, dx \\ &= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\frac{L_d - L_s + x}{T}\frac{S}{\lambda} g^{S+1}(x) + h(L_d - L_s)\frac{L_d - L_s + x}{T} g^S(x) \\ &\quad + (h+l)\frac{T^2 - (T - L_d + L_s - x)^2}{2T} g^S(x)\, dx \\ &\quad + \int_{L_s-(L_d-T)^+}^{L_s} h\frac{S}{\lambda} g^{S+1}(x) + h(L_d - L_s) g^S(x) + (h+l)\frac{T}{2} g^S(x)\, dx \end{aligned}$$

$$
\begin{aligned}
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\frac{L_d-L_s}{T}\frac{S}{\lambda}g^{S+1}(x) + h\frac{S(S+1)}{T\lambda^2}g^{S+2}(x) + h\frac{(L_d-L_s)^2}{T}g^S(x) \\
&\quad + h\frac{L_d-L_s}{T}\frac{S}{\lambda}g^{S+1}(x) + (h+l)\frac{T^2-(T-L_d+L_s)^2}{2T}g^S(x)\,dx \\
&\quad + (h+l)\frac{2(T-L_d+L_s)}{2T}\frac{S}{\lambda}g^{S+1}(x) - (h+l)\frac{S(S+1)}{T\lambda^2}g^{S+2}(x) \\
&\quad + \int_{L_s-(L_d-T)^+}^{L_s} h\frac{S}{\lambda}g^{S+1}(x) + h(L_d-L_s)g^S(x) + (h+l)\frac{T}{2}g^S(x)\,dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\left(\frac{2(L_d-L_s)^2}{2T} + \frac{T^2-(T-L_d+L_s)^2}{2T}\right)g^S(x) \\
&\quad + h\left(\frac{L_d-L_s}{T} + \frac{L_d-L_s}{T} + \frac{T-L_d+L_s}{T}\right)\frac{S}{\lambda}g^{S+1}(x) \\
&\quad + h\frac{S(S+1)}{2T\lambda^2}g^{S+2}(x) + l\frac{T^2-(T-L_d+L_s)^2}{2T}g^S(x) \\
&\quad + l\frac{(T-L_d+L_s)S}{T\lambda}g^{+1}S(x) - l\frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\,dx \\
&\quad + \int_{L_s-(L_d-T)^+}^{L_s} h\left(\left(L_d-L_s+\frac{T}{2}\right)g^S(x) + \frac{S}{\lambda}g^{S+1}(x)\right) + l\frac{T}{2}g^S(x)\,dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\left(\left(\frac{2L_d^2-4L_sL_d+2L_s^2+T^2-T^2-L_d^2-L_s^2+2TL_d-2TL_s}{2T}\right.\right. \\
&\quad \left.\left. + \frac{2L_sL_d}{2T}\right)g^S(x) + \frac{(L_d-L_s+T)S}{T\lambda}g^{S+1}(x) + \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\right) \\
&\quad + l\left(\frac{T^2-T^2-L_d^2-L_s^2+2TL_d-2TL_s+2L_sL_d}{2T}g^S(x)\right. \\
&\quad \left. + \frac{(T-L_d+L_s)S}{T\lambda}g^{S+1}(x) - \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\right)\,dx \\
&\quad + \int_{L_s-(L_d-T)^+}^{L_s} h\left(\left(L_d-L_s+\frac{T}{2}\right)g^S(x) + \frac{S}{\lambda}g^{S+1}(x)\right) + l\frac{T}{2}g^S(x)\,dx \\
&= \int_{L_s-L_d}^{L_s-(L_d-T)^+} h\left(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}g^S(x) + \frac{(L_d-L_s+T)S}{T\lambda}g^{S+1}(x)\right. \\
&\quad \left. + \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\right) + l\left(\frac{-(L_d-L_s)^2+2T(L_d-L_s)}{2T}g^S(x)\right. \\
&\quad \left. + \frac{(T-L_d+L_s)S}{T\lambda}g^{S+1}(x) - \frac{S(S+1)}{2T\lambda^2}g^{S+2}(x)\right)\,dx \\
&\quad + \int_{L_s-(L_d-T)^+}^{L_s} h\left(\left(L_d-L_s+\frac{T}{2}\right)g^S(x) + \frac{S}{\lambda}g^{S+1}(x)\right) + l\frac{T}{2}g^S(x)\,dx \\
&= h\left(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}\Big(G^S(L_s-(L_d-T)^+) - G^S(L_s-L_d)\Big)\right.
\end{aligned}
$$

$$
\begin{aligned}
&+ \frac{(L_d - L_s + T)S}{T\lambda}\Big(G^{S+1}(L_s - (L_d - T)^+) - G^{S+1}(L_s - L_d)\Big) \\
&+ \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - (L_d - T)^+) - G^{S+2}(L_s - L_d)\Big)\Bigg) \\
&+ l\Bigg(\frac{2T(L_d - L_s) - (L_d - L_s)^2}{2T}\Big(G^{S}(L_s - (L_d - T)^+) - G^{S}(L_s - L_d)\Big) \\
&+ \frac{(T - L_d + L_s)S}{T\lambda}\Big(G^{S+1}(L_s - (L_d - T)^+) - G^{S+1}(L_s - L_d)\Big) \\
&- \frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s - (L_d - T)^+) - G^{S+2}(L_s - L_d)\Big)\Bigg) \\
&+ h\Bigg(\bigg(L_d - L_s + \frac{T}{2}\bigg)\Big(G^{S}(L_s) - G^{S}(L_s - (L_d - T)^+)\Big) \\
&+ \frac{S}{\lambda}\Big(G^{S+1}(L_s) - G^{S+1}(L_s - (L_d - T)^+)\Big)\Bigg) \\
&+ l\frac{T}{2}\Big(G^{S}(L_s) - G^{S}(L_s - (L_d - T)^+)\Big) \qquad (4.87)
\end{aligned}
$$

The expected inventory costs for Situation F are given as:

$$
\begin{aligned}
&E[C_F(\Omega(S),V)]\\
&=\begin{cases}
h\Big(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}\Big(G^S(L_s)-G^S(L_s-L_d)\Big)\\
+\frac{(L_d-L_s+T)S}{T\lambda}\Big(G^{S+1}(L_s)-G^{S+1}(L_s-L_d)\Big)\\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s)-G^{S+2}(L_s-L_d)\Big)\Big)\\
+l\Big(\frac{2T(L_d-L_s)-(L_d-L_s)^2}{2T}\Big(G^S(L_s)-G^S(L_s-L_d)\Big)\\
+\frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s)-G^{S+1}(L_s-L_d)\Big)\\
-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s)-G^{S+2}(L_s-L_d)\Big)\Big) & \text{if } L_d\le T,\\
h\Big(\frac{(L_d-L_s)^2+2T(L_d-L_s)}{2T}\Big(G^S(L_s-L_d+T)-G^S(L_s-L_d)\Big)\\
+\frac{(L_d-L_s+T)S}{T\lambda}\Big(G^{S+1}(L_s-L_d+T)-G^{S+1}(L_s-L_d)\Big)\\
+\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-L_d+T)-G^{S+2}(L_s-L_d)\Big)\\
+(L_d-L_s+\frac{T}{2})\Big(G^S(L_s)-G^S(L_s-L_d+T)\Big)\\
+\frac{S}{\lambda}\Big(G^{S+1}(L_s)-G^{S+1}(L_s-L_d+T)\Big)\Big)\\
+l\Big(\frac{2T(L_d-L_s)-(L_d-L_s)^2}{2T}\Big(G^S(L_s-L_d+T)-G^S(L_s-L_d)\Big)\\
+\frac{(T-L_d+L_s)S}{T\lambda}\Big(G^{S+1}(L_s-L_d+T)-G^{S+1}(L_s-L_d)\Big)\\
-\frac{S(S+1)}{2T\lambda^2}\Big(G^{S+2}(L_s-L_d+T)-G^{S+2}(L_s-L_d)\Big)\\
+\frac{T}{2}\Big(G^S(L_s)-G^S(L_s-L_d+T)\Big)\Big) & \text{otherwise.}
\end{cases}
\end{aligned}
\tag{4.88}
$$

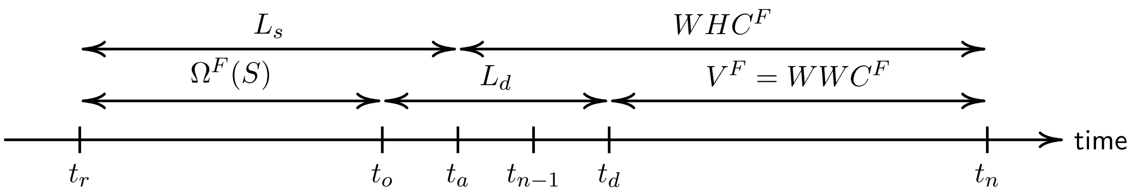

Figure 4.22: Timeline with the important time instances for Situation F when $S > 0$

Situation G:

Since the considered unit is available after it is demanded, early deliveries are not possible.

$$\begin{aligned} F_G(x,y) &= P\Big(\Omega(S) \le x, V \le y, t_o \le t_a, t_o < t_d \le t_a < t_n, t_s = t_n\Big) \\ &= P\Big(\Omega(S) \le x, V \le y, t_r + \Omega(S) \le t_r + L_s, \\ &\quad t_r + \Omega(S) < t_r + \Omega(S) + L_d \le t_r + L_s < t_n, t_s = t_n\Big) \\ &= P\Big(\Omega(S) \le x, V \le y, \Omega(S) \le L_s, \Omega(S) < \Omega(S) + L_d \le L_s\Big) \end{aligned} \quad (4.89)$$

For $x \le L_s - L_d$ we get

$$\begin{aligned} F_G(x,y) &= P\Big(\Omega(S) \le x, V \le y, \Omega(S) \le L_s, \Omega(S) < \Omega(S) + L_d \le L_s\Big) \\ &= G^S(x)U(y). \end{aligned} \quad (4.90)$$

Thus, the joint density is given as

$$f_G(x,y) = \begin{cases} g^S(x)u(y) & x \le L_s - L_d, y \le T \\ 0 & \text{otherwise.} \end{cases} \quad (4.91)$$

$$\begin{aligned} &E[C_G(\Omega(S), V)] \\ &= \int_0^\infty \int_0^T C_G(x,y) f_G(x,y)\, dy\, dx \\ &= \int_0^{L_s-L_d} \int_0^T \Big(hy + l(y + L_s - L_d - x)\Big) g^S(x)u(y)\, dy\, dx \\ &= \int_0^{L_s-L_d} (h+l)\frac{T}{2} g^S(x) + l(L_s - L_d) g^S(x) - l\frac{S}{\lambda} g^{S+1}(x)\, dy\, dx \\ &= h\frac{T}{2} G^S(L_s - L_d) + l\left(\left(L_s - L_d + \frac{T}{2}\right) G^S(L_s - L_d) - \frac{S}{\lambda} G^{S+1}(L_s - L_d)\right) \end{aligned} \quad (4.92)$$

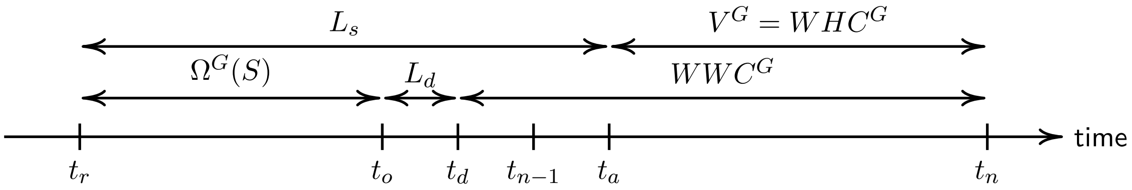

Figure 4.23: Timeline with the important time instances for Situation G when $S > 0$

4.6.3 Derivations of the Expected Inventory Costs when $S \leq 0$

The expected inventory costs when $S \leq 0$ comply with the expected warehouse holding costs in Situation G when $S > 0$ because in both cases the unit will always be first demanded and then it is available. If $S = 0$, the unit is ordered at the same time point where the facility order arrives, whereas $S < 0$ means that we order the considered unit after the next $|S|$ facility orders are placed. The expected inventory costs are derived below. Early deliveries are not possible. The time points when the unit is ordered, demand or available can be described as $t_o = t_r - \Omega(|S|) < t_d = t_r - \Omega(|S|) + L_d \leq t_a = t_r + L_s$.

$$\begin{aligned}\tilde{F}(x, y) &= P\Big(\Omega(|S|) \leq x, V \leq y, t_o < t_d \leq t_a\Big) \\ &= P\Big(\Omega(|S|) \leq x, V \leq y, t_r - \Omega(|S|) < t_r - \Omega(|S|) + L_d \leq t_r + L_s\Big) \\ &= P\Big(\Omega(|S|) \leq x, V \leq y, 0 < L_d \leq \Omega(|S|) + L_s\Big) \end{aligned} \tag{4.93}$$

For $x \leq \infty$ we get

$$\begin{aligned}\tilde{F}(x, y) &= P\Big(\Omega(|S|) \leq x, V \leq y, 0 < L_d \leq \Omega(|S|) + L_s\Big) \\ &= P\Big(\Omega(|S|) \leq x, V \leq y\Big) \\ &= G^{|S|}(x)U(y) \end{aligned} \tag{4.94}$$

Thus, the joint probability density function is given as

$$\tilde{f}(x, y) = \begin{cases} g^{|S|}(x)u(y) & x \leq \infty, y \leq T \\ 0 & \text{otherwise.} \end{cases} \tag{4.95}$$

During V time units the warehouse has to keep stock on hand and during an expected time interval of $V + L_s - L_d + \Omega(|S|)$ late-delivery costs occur, why we get the following equations.

$$\begin{aligned}E[\tilde{C}(\Omega(S), V)] &= \int_0^\infty \int_0^T \tilde{C}(x, y)\tilde{f}(x, y)\, dy\, dx \\ &= \int_0^\infty \int_0^T \Big(hy + l(y + L_s - L_d + x)\Big)g^{|S|}(x)u(y)\, dy\, dx \\ &= \int_0^\infty \int_0^T \Big((h + l)y + l(L_s - L_d) + lx\Big)g^{|S|}(x)u(y)\, dy\, dx \\ &= \int_0^\infty (h + l)\frac{T}{2}g^{|S|}(x) + l(L_s - L_d)g^{|S|}(x) + l\frac{|S|}{\lambda}g^{S+1}(x)\, dy\, dx \end{aligned}$$

$$= h\frac{T}{2}G^{|S|}(\infty) + l\Bigg(\bigg(L_s - L_d + \frac{T}{2}\bigg)G^{|S|}(\infty) + \frac{|S|}{\lambda}G^{S+1}(\infty)\Bigg)$$

$$= h\frac{T}{2} + l\bigg(\frac{|S|}{\lambda} + L_s - L_d + \frac{T}{2}\bigg) \tag{4.96}$$

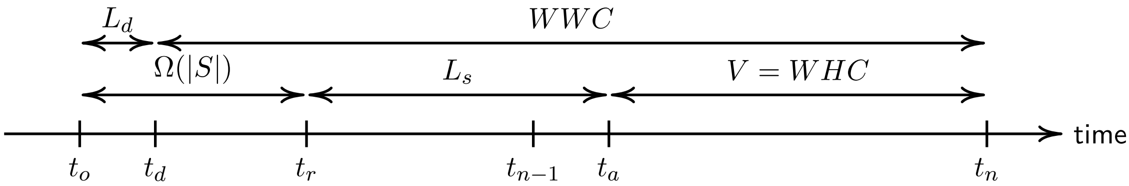

Figure 4.24: Timeline with the important time instances when $S \leq 0$

4.6.4 Total Cost of the System without Flexible Deliveries

For our numerical study we need to calculate the shipment costs and inventory costs when ADI is available but flexible deliveries are not allowed. In that case, the shipment can be obtained in a similar way as shown in Stenius et al. (2018). Consider that advance orders can only be shipped after they are due. When computing the expected inventory costs, there is no decision about early shipments. Therefore, we do not have to distinguish between all 7 situations as before, we only distinguish between Situation A ($t_a \leq t_o$) and Situation G ($t_o < t_a$). This yields inventory costs at the warehouse of

$$E[\bar{C}_A(\Omega(S))] = h\Bigg(\bigg(L_d - L_s + \frac{T}{2}\bigg)\Big(1 - G^S(L_s - L_d)\Big) + \frac{S}{\lambda}\Big(1 - G^{S+1}(L_s - L_d)\Big)\Bigg) + l\frac{T}{2}\Big(1 - G^S(L_s - L_d)\Big) \tag{4.97}$$

$$E[\bar{C}_G(\Omega(S))] = h\frac{T}{2}G^S(L_s - L_d) + l\bigg(\Big(L_s - L_d + \frac{T}{2}\Big)G^S(L_s - L_d) - \frac{S}{\lambda}G^{S+1}(L_s - L_d)\bigg) \tag{4.98}$$

Chapter

5 Optimal Time-based Shipment Policy with Flexible Deliveries

The following chapter is based on Ralfs et al. (2024). Building upon the insights from the previous chapter, this chapter again focuses on the joint consideration of shipment consolidation and ADI. However, this chapter differs in several key aspects, as outlined in Section 5.1. We investigate a periodic review single-echelon inventory system that satisfies random orders from a production facility. The facility does not expect the orders to be delivered immediately, but it offers ADI in form of a demand lead time. In contrast to the previous chapter, we do not assume heuristic shipment quantities but instead optimize these quantities given a time-based dispatch schedule. In Section 5.2, we present the MDP, including the state, action, transition functions and period cost. On this basis, Section 5.3 describes the DRL-based solution algorithm, which relies on state value approximation, along with an explanation of how the deep neural network is trained. The results, as discussed in Section 5.4, yield that the DRL policy achieves optimal decisions in small scale instances. Additionally, we observe general structural properties of the near-optimal outbound shipment quantities and propose an approximate threshold (ATH) policy, which performs very well in most of the large-scale instances. We further benchmark the DRL policy and the ATH policy against three simple policies. Section 5.5 provides a conclusion of this chapter, summarizing the key findings and insights.

5.1 Problem Formulation

We investigate a periodic review inventory model in which a single warehouse satisfies stochastic orders from a production facility for a single item at the beginning of a shipment period. The production facility applies a preorder strategy, meaning that every order does not have to be satisfied directly but comes along with a due date. This strategy is represented by a demand lead time L_d, which specifies the number of periods between the order placement by the facility and the due date. For instance, an order placed at the end of period i is due at the beginning of period $i + L_d + 1$. A preorder does not have to be satisfied directly until it is due. Therefore, the inventory manager distinguishes between *orders* (orders which are not due) and *demands* (orders which are due). For simplicity, we assume that the production facility always orders with an identical and constant demand lead time L_d and that order cancellations are negligible due to high cancellation costs. We do not consider an explicit replenishment policy at the warehouse, and we assume that there is always sufficient stock at the warehouse so that we can focus purely on the optimal outbound shipment policy.

The warehouse applies a time-based shipment dispatch policy, meaning that a shipment is dispatched to the production facility every T periods. The time T between two shipment periods is called the *shipment interval*. In contrast to approaches in the literature on time-based shipment consolidation, the warehouse does not ship all accumulated demands accumulated since the last dispatch. Instead, we assume that the inventory manager decides about the shipment quantity at the beginning of a shipment period, taking into account the information about the number of demands and orders. Like Wang and Toktay (2008), the warehouse applies the concept of flexible deliveries. Therefore, orders can be shipped before the corresponding due date. Instead of making individual binary decisions for each ordered unit, the inventory manager determines the total number of units to be dispatched at the beginning of a shipment period t, denoted by x_t. This process results in a single integer-based decision about all demands and orders, with units being dispatched based on the FCFS principle. In summary, when the inventory manager chooses a shipment quantity less than the number of demands, only the demands are shipped, whereas a shipment quantity larger than the number of demands leads to a shipment of all demands and additional orders.

Let us enumerate the decision periods as $t \in \mathcal{T}$, where $\mathcal{T}$ defines the set of decision periods $\mathcal{T} = \{0, 1T, 2T, \dots\}$ with T as the shipment interval, and we define the sequence of events as follows:

(*i*) Review of the system state at the beginning of the shipment period t

(*ii*) Decision about the outbound shipment quantity to the production facility at the beginning of the shipment period t

(*iii*) Arrival of T stochastic orders for periods $t, t+1, \ldots, t+T-1$ with identical demand lead times at the end of the shipment period t

(*iv*) Incursion of cost at the end of the shipment period t

Now, we focus on the cost incurred within this inventory system. A 3PL manages the movement of goods from the warehouse to the production facility to be able to act flexibly. The terms of the contract with the 3PL are set as follows: For each shipment period t, a prearranged transportation capacity of Cap units is reserved, referred to as *primary transportation option*. If the actual shipment quantity in shipment period t, x_t, exceeds this reserved capacity Cap, the manager ships $(x_t - Cap)^+$ units with an *alternative transportation option*. We assume linear shipment costs c_1 per reserved transportation capacity for the primary transportation option. These costs are similar in each shipment period and correspond to $c_1 \cdot Cap$. These reservation costs are not decision relevant and, therefore, are not included in our analysis; however, it is simple to add them to managerial proposals in a numerical study. Each unit shipped by the alternative transportation option incurs linear shipment costs c_2. The managers' decision directly influences this part of the shipment costs. We note that $c_1 \leq c_2$.

In addition to shipment costs, two other types of costs are incurred. First, if orders are dispatched to the production facility before the due date, early-delivery costs e are incurred per unit and period. This term represents the additional inventory holding costs at the production facility, where space is often limited. Second, if the warehouse manager decides not to fulfill demands, late-delivery costs l are incurred per unit and period. This reflects the cost of higher safety stocks at the production facility to hedge against stockouts.

Our goal is to determine the optimal outbound shipment quantities for a given shipment interval T and a given reserved transportation capacity Cap under random orders, which minimizes the expected total cost per period comprising early-delivery, late-delivery, and shipment costs throughout the planning horizon. An overview of the problem is presented in Figure 5.1.

Shipping all the demands and orders to the production facility at the beginning of a shipment period may lead to high early-delivery and shipment costs if the reserved capacity Cap is exceeded. Shipping at maximum Cap units minimizes shipment costs; however, this strategy may lead to high late-delivery costs. Therefore, the warehouse manager

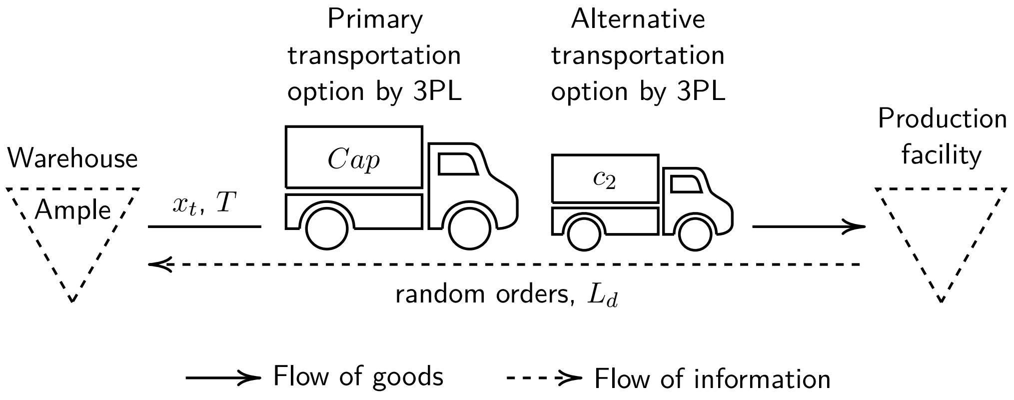

Figure 5.1: Flow of information and goods in the considered inventory system

must determine the shipment quantity in such a way that the right balance of these types of costs is found.

5.2 Markov Decision Process

We formalize the described problem by using a MDP, with the components detailed as follows:

Decision epoch. Let us consider an infinite planning horizon. The inventory manager makes decisions at the beginning of each shipment period $t \in \mathcal{T}$.

Pre-decision state. The state at the beginning of shipment period t, before the decision takes place, is an $(L_d + 1)$-dimensional vector, $\bar{S}_t = \left(w_t, y_t^1, y_t^2, \ldots, y_t^{L_d}\right)$. The component w_t denotes the demands already due in shipment period t, and y_t^n signifies the orders due in period $t + n$.

Decision. Let $\mathcal{X}_t$ be the set of potential decisions at the beginning of shipment period t and let x_t indicate the chosen decision at that time. Given that the inventory manager decides the quantity of units to dispatch, $\mathcal{X}_t$ includes all integers from zero up to the total of demands and orders. Thus, x_t belongs to the set $\{0, 1, \ldots, w_t + \sum_{i=1}^{L_d} y_t^i\}$.

Post-decision state. The state after making a decision, represented as $\bar{S}_t^x$, encompasses the information available once the decision x_t has been executed but prior to the arrival of any exogenous information. In terms of dimensionality, $\bar{S}_t^x$ is identical to the pre-decision state $\bar{S}_t$. Each dimension is updated in the following manner:

$$\tilde{w}_t = (w_t - x_t)^+ \tag{5.1}$$

$$\tilde{y}_t^j = \min\left(\left(w_t + \sum_{i=1}^{j} y_t^i - x_t\right)^+, y_t^j\right) \quad \forall\, j \in \{1, 2, \ldots, L_d\} \tag{5.2}$$

Given that orders are dispatched based on the FCFS principle, we subtract the dispatch decision from the pre-decision state vector in sequence. We begin with the demands, as shown in Equation (5.1). We then proceed with the orders from the lowest due date to the highest due date, as shown in Equation (5.2). The post-decision state can then be written as $\bar{S}_t^x = \left(\tilde{w}_t, \tilde{y}_t^1, \tilde{y}_t^2, \ldots, \tilde{y}_t^{L_d}\right)$.

Exogenous information. We define a vector Φ_t with T dimensions that contains the orders from the production facility between two decision points. Since orders are received on a periodic basis, we denote $\Phi_t = \left(D_t^0, D_t^1, \ldots, D_t^{T-1}\right)$, where D_t^n is a random variable reflecting the number of ordered units at the end of period $t+n$ and due in period $t+n+L_d+1$, with $t \in \mathcal{T}$. The realization of Φ_t is denoted by $\phi_t = \left(d_t^0, d_t^1, \ldots, d_t^{T-1}\right)$.

Transition functions and costs. Once the exogenous information variable Φ_t materializes, the system transitions to the subsequent pre-decision state $\bar{S}_{t+T}$, and associated costs incur. The transition to the next pre-decision state $\bar{S}_{t+T} = \left(w_{t+T}, y_{t+T}^1, y_{t+T}^2, \ldots, y_{t+T}^{L_d}\right)$ is described in Equation (5.3) and Equation (5.4).

$$w_{t+T} = \tilde{w}_t + \sum_{i=1}^{\min(L_d,T)} \tilde{y}_t^i + \sum_{i=0}^{(T-L_d)^+-1} d_t^i \tag{5.3}$$

$$y_{t+T}^j = \begin{cases} d_t^{j+T-L_d-1} & \forall\, j \in \{1, 2 \ldots, L_d\} & \text{if } L_d \leq T \\ \tilde{y}_t^{j+T} & \forall\, j \in \{1, 2, \ldots, L_d - T\} & \\ d_t^{j-L_d+1} & \forall\, j \in \{L_d - T + 1, L_d - T + 2, \ldots, L_d\} & \text{if } L_d > T \end{cases} \tag{5.4}$$

Since the transition depends on the length of the demand lead time and the length of the shipment interval, we distinguish between two scenarios: $L_d \leq T$ and $L_d > T$.

The demand transition in period $t+T$ can be presented consistently, including both cases. The first term in Equation Equation (5.3) represents that all the demands remaining unfulfilled after the shipment decision, i.e., $\tilde{w}_t$, must be added to the demands in period $t+T$, w_{t+T}. Additionally, orders that have not been shipped and became due before $t+T$ are incorporated into w_{t+T}. The third term in Equation Equation (5.3) demonstrates that the orders received at the end of shipment period t are attributed to w_{t+T} if the demand lead time is less than or equal to the length of the shipment interval, meaning that their due date is before or at the subsequent shipment period $t+T$. Consequently, there is no chance for these orders to be dispatched before their due dates.

To illustrate the transition of orders with the corresponding due date j, y^j_{t+T} $\forall\ j \in \{1, 2, \ldots, L_d\}$, we need to separate the transition. If $L_d \leq T$, all known orders that remain after the shipment decision mutate to demands. Therefore, only new orders received at the end of shipment period t are assigned to y^j_{t+T} $\forall\ j \in \{1, 2, \ldots, L_d\}$ with a due date in period $t+T+j$. However, if $L_d > T$, known orders that have not been shipped at t and whose due date is after $t+T$ continue to be orders. Additionally, all new orders received at the end of shipment period t are due after $t+T$.

In the sequel, we derive the expressions for the expected total cost per period, when a shipment decision is made every T periods. Provided that the shipment decision is set at the beginning of a shipment period $t \in \mathcal{T}$, we can easily compute the additional shipment costs when using the alternative transportation option.

$$c_t^{shipment}(x_t) = c_2(x_t - Cap)^+ \tag{5.5}$$

Furthermore, all units shipped before their due date incur early-delivery costs.

$$c_t^{early}(x_t) = e\sum_{i=1}^{T}(y_t^i - \tilde{y}_t^i)i \tag{5.6}$$

Moreover, the expected late-delivery costs are defined as follows:

$$\begin{aligned} c_t^{late}(x_t) = l\Bigg(& T\tilde{w}_t + \sum_{i=1}^{\min(L_d,T-1)} \tilde{y}_t^i(T-i) \\ &+ \mathbb{E}\Bigg[\sum_{i=0}^{\min(L_d-T,L_d-1)} D_t^i(T-L_d-i-1)\Bigg]\Bigg) \end{aligned} \tag{5.7}$$

In summary, the expected total cost per period when the system is in state $\bar{S}_t$ and making decision x_t can be written as follows:

$$c(\bar{S}_t, x_t) = \frac{1}{T}\left(c_t^{shipment}(x_t) + c_t^{early}(x_t) + c_t^{late}(x_t)\right) \tag{5.8}$$

Let Π represent the set of feasible policies. A policy $\pi \in \Pi$ is a function that maps a pre-decision state $\bar{S}_t$ to a decision x_t. The objective is to identify the optimal policy π^* that minimizes the expected total discounted cost C_t, starting from the initial state $\bar{S}_0$.

$$\pi^* = \underset{\pi \in \Pi}{\operatorname{argmin}} \left\{ \sum_{t \in \{0, 1T, 2T, \dots\}} \gamma c(\bar{S}_t, \pi(\bar{S}_t)|\bar{S}_0) \right\}, \tag{5.9}$$

where $\gamma \in (0, 1]$ is the discount factor.

In theory, the value function of a state $\bar{S}_t$, denoted by $V(\bar{S}_t)$, which is the minimum cumulative discounted cost when starting from state $\bar{S}_t$, can be recursively determined by using Bellman's equation:

$$V(\bar{S}_t) = \min_{x_t \in \mathcal{X}_t} \{c(\bar{S}_t, x_t)T + \gamma \sum_{\bar{S}' \in \mathcal{S}_{t+T}} Pr(\bar{S}_{t+T} = \bar{S}'|\bar{S}_t, x_t)V(\bar{S}')\}, \tag{5.10}$$

where $\mathcal{S}_{t+T}$ denotes the set of all possible pre-decision states in period $t + T$.

Once the state values $V(\bar{S}_t)$ for all states in Equation (5.10) are determined, the optimal decision can then be extracted by using:

$$x_t^* = \underset{t \in \mathcal{T}}{\operatorname{argmin}} \{c(\bar{S}_t, x_t)T + \gamma \sum_{\bar{S}' \in \mathcal{S}_{t+T}} Pr(\bar{S}_{t+T} = \bar{S}'|\bar{S}_t, x_t)V(\bar{S}')\} \tag{5.11}$$

While the Bellman optimality equations represented in Equation (5.10) and Equation (5.11) provide an exact framework for computing the optimal policy, their applicability is generally limited to small-scale problem instances due to the well-known *curse of dimensionality* (Powell, 2022). In the following section, we introduce an algorithm designed to approximate the optimal policy in a computationally efficient manner.

5.3 Deep Reinforcement Learning Policy

In this section, we present our DRL-based solution algorithm to obtain an outbound shipment policy, denoted by π^{DRL}. In Section 5.4, we demonstrate that the approach under discussion yields near-optimal decisions in small-scale instances and performs well in large-scale instances. In the following, we first outline the fundamental concepts before delving into the specifics of training the deep neural network.

5.3.1 General Concepts

Applying the state value approximation provides a way to circumvent the limitations of size and dimensionality in the decision space of policy-based reinforcement learning approaches. In contrast to policy-based methods, where decision-making is restricted by the number of output nodes in the policy network, the state value approximation enables a more general optimization problem, as stated in Equation (5.11). However, even when the pre-decision state values in Equation (5.10) are accurately approximated, extensive forward simulation is typically necessary to calculate the expected cost-to-go for making a decision by using Equation (5.11). This becomes an issue when the transition space is large, which is common in real-world problems. Therefore, we opt to approximate the value of the post-decision state $\bar{S}_t^x$, as described in Powell (2022), to partially circumvent this issue.

Given that the post-decision state captures the condition of the system after a decision has been made but prior to the arrival of any exogenous information, it follows that the immediate deterministic costs, denoted by $c^{det}(\bar{S}_t, x_t)$, include all costs that do not depend on the realization of the vector Φ_t. Accordingly, $c^{det}(\bar{S}_t, x_t)$ is calculated as follows:

$$c^{det}(\bar{S}_t, x_t) = c_2(x_t - Cap)^+ + e\sum_{i=1}^{T}(y_t^i - \tilde{y}_t^i)i + l\left(T\tilde{w}_t + \sum_{i=1}^{\min(L_d, T-1)} \tilde{y}_t^i(T-i)\right) \quad (5.12)$$

The optimization problem now becomes:

$$x_t^* = \underset{x_t \in \mathcal{X}_t}{\operatorname{argmin}}\{c^{det}(\bar{S}_t, x_t) + \gamma V_x(\bar{S}_t^x)\} \quad (5.13)$$

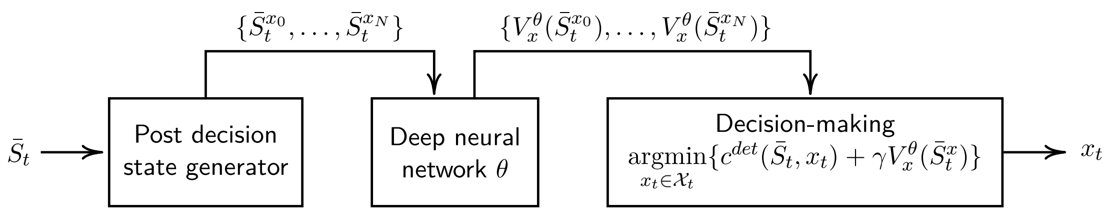

Figure 5.2: Decision-making process of the DRL policy

where $V_x(\bar{S}_t^x)$ is the value function of the post-decision state $\bar{S}_t^x$, which is the minimum cumulative discounted cost when starting from the post-decision state $\bar{S}_t^x$.

A general overview of the process, from being in a state $\bar{S}_t$ and using π^{DRL} to finally making a decision x_t, is illustrated in Figure 5.2. Starting from the pre-decision state $\bar{S}_t$, all possible post-decision states are generated by iterating through each decision from 0 to N, where N is the total number of demands and orders. Subsequently, a deep neural network, represented by θ, assigns a value to each generated post-decision state. The final decision is determined by choosing the option that minimizes the sum of the immediate deterministic costs, $c^{det}(\bar{S}_t, x_t)$, and the value of the associated post-decision state. Specifically:

$$\pi^{DRL}(\bar{S}_t) = \underset{x_t \in \mathcal{X}_t}{\operatorname{argmin}}\{c^{det}(\bar{S}_t, x_t) + \gamma V_x^\theta(\bar{S}_t^x)\} \tag{5.14}$$

Let us note that no forward simulation was performed, as all terms in Equation (5.14) can be calculated exactly.

In problems with a relatively small decision space, this modification enables near-instantaneous decision-making since no forward simulation is needed to estimate the cost-to-go of making a decision x_t; the post-decision state $\bar{S}_t^x$ already captures this information. This not only speeds up the decision process when the policy is already in place but also facilitates rapid roll-out during training to aid post-decision state value approximation. This approach is also applicable to problems with much larger decision spaces; however, in those cases, an efficient method must be found to solve the *deterministic* optimization problem presented in Equation (5.13).

However, the post-decision state value function in Equation (5.14) needs to be accurately approximated. To this end, we use a simple fully connected neural network and train it to map post-decision states to their corresponding values. We elaborate on this in the subsequent section.

5.3.2 Training the Deep Neural Network

In this approach, we need a function to map a multidimensional post-decision state vector into a scalar value representation. To achieve this, we chose a straightforward fully connected network. This network has $L_d + 1$ input nodes and consists of three hidden layers, each with 128 nodes, and we use the rectified linear unit (ReLU) as the activation function. The final output layer utilizes linear activation and has a single output node.

The deep neural network is trained by using the following procedure. Initially, the parameters of the network, denoted by θ, are subject to random initialization. We then instantiate a random state $\bar{S}_t$ by randomly initializing the elements of its state vector. Subsequent decisions are made according to Equation (5.13) by using the current parameter set θ to compute the value of the post-decision state. At the outset of training, the values produced by the network are not expected to provide accurate approximations of the post-decision state due to the purely random initialization of θ. After making the decision x_t, we arrive at the post-decision state $\bar{S}_t^x$ before the system is simulated forward to the subsequent state $\bar{S}_{t+T}$ using a *single* realization of the random vector Φ_t. Given this realization ϕ_t, we can calculate the associated cost as follows:

$$c^{rand}(\bar{S}_t, x_t) = \sum_{i=0}^{\min(L_d - T, L_d - 1)} d_t^i(T - L_d - i - 1) \tag{5.15}$$

In the subsequent pre-decision state $\bar{S}_{t+T}$, a decision is once again selected by using the previously described procedure. This yields the next deterministic cost $c^{det}(\bar{S}_{t+T}, x_{t+T})$ as well as the following post-decision state $\bar{S}_{t+T}^x$. We then store the following quadruples in a double-ended queue, commonly referred to as the *experience buffer*, as outlined in Mnih et al. (2015):

$$\left(\bar{S}_t^x, c^{rand}(\bar{S}_t, x_t), c^{det}(\bar{S}_{t+T}, x_{t+T}), \bar{S}_{t+T}^x\right) \tag{5.16}$$

This entire procedure continues until an adequate number of experiences is collected in the buffer. Once this threshold is reached (in the context of our study, 64 experience quadruples are in the buffer), the optimization of the neural network θ begins. A random batch of experience is sampled from the buffer, and the loss function from experience is constructed as follows:

$$\mathcal{L} = [V_x^{\theta}(\bar{S}_t^x) - (c^{rand}(\bar{S}_t, x_t) + c^{det}(\bar{S}_{t+T} + x_{t+T}) + V_x^{\theta^{target}}(\bar{S}_{t+T}^x))]^2, \tag{5.17}$$

where θ^{target} serves as the *target network* (Mnih et al., 2015) and is structurally identical to θ. This target network is a delayed duplicate of θ and receives updates at fixed intervals of every I iterations. Subsequently, the network θ undergoes an update by using the well-known Adam optimizer (Kingma and Ba, 2014) to minimize the loss function defined in Equation (5.17). We employ the ϵ-greedy approach, in which there is an ϵ probability of selecting a random, feasible decision to promote exploration of the state space.

In our study, we configured the experience buffer to have a length of 1,000 experiences, set the target network update interval I at 100 iterations, and chose a batch size of 64. The initial value of ϵ is set to 1, and the value undergoes exponential decay at a rate of 0.999 in each iteration. Varying these hyperparameters within a similar range does not significantly impact the training results, indicating that the model is relatively robust to changes in these settings. To capitalize on the benefits of estimating post-decision state values for faster decision-making, we conduct a brief roll-out (without ϵ-greedy) every 1,000 iterations. During this process, we use the current network parameters to assess the current out-of-sample performance. Each roll-out lasts for 10,000 epochs. If the current network outperforms previous best-performing networks in these rollouts (i.e., achieves a lower total cost), its parameters are retained as the best-found network parameters, denoted by θ^{best}. In addition, if no improvement was found in the last ten roll-outs, we then reset the parameter ϵ to one to diversify the search. The whole training process stops after 100,000 epochs.

At the end of the training process, we obtain the best-found neural network parameters θ^{best}, which can then be used in Equation (5.13) to make decisions on unseen instances in the numerical experiment.

5.4 Numerical Study

To evaluate our proposed DRL approach, we divide the numerical experiments into three parts. First, we compute the expected total cost per period when using the DRL policy for a selection of small-scale scenarios and compare them with the expected total cost per period when using the optimal policy, computed with VI. We can observe several structural properties of the optimal policy, which we use to define an ATH policy. In addition, we present simple but reasonable heuristics without any thresholds. Finally, we turn our focus to a set of large-scale instances characterized by extended demand lead times, longer shipment intervals, and a higher order rate. Due to the computational limitations of VI for such large instances, we benchmark the DRL policy π^{DRL} against all heuristics. Finally, we draw managerial insights; in particular, we focus on the value of ADI and transportation capacity planning.

We compute the average total cost per period of our policies for each instance via simulation. The length of each simulation replication is 10,000 periods, while excluding a warm-up period of 1,000 periods in each replication. A minimum of ten replications is performed. Afterward, we use sequential sampling and stop if the half-width of the 95 % confidence interval of the average total cost per period is less than 1 % of the current average total cost per period of the considered instance.

5.4.1 Small-Scale Instances

The settings for small-scale scenarios are outlined in Table 5.1. For each period within the planning horizon, the facility's order quantities are determined randomly based on a uniform distribution, $\mathcal{U}(d_{\text{min}}, d_{\text{max}})$, where both endpoints are inclusive. The reserved transportation capacity Cap is calculated by taking the product of the capacity ratio r_{Cap} and the length of the shipment interval T and then multiplying it by the mean of the aforementioned period order quantities, $\frac{d_{\text{min}}+d_{\text{max}}}{2}$. The capacity ratio r_{Cap} reflects the planner's strategy for capacity reservation. If r_{Cap} is less than 1, Cap is less than the expected number of facility orders between two shipment periods, and vice versa.

The set of parameter values detailed in Table 5.1 yields 216 unique instances. The performance gap (i.e., the relative difference in the expected total cost per period)

Table 5.1: Problem parameters for small instances

Parameters		Values
Shipment interval	T	{2, 4}
Demand lead time	L_d	{1, 2}
Capacity ratio	r_{Cap}	{0.5, 1, 2}
Shipment cost	c_2	{1, 10, 50}
Early-delivery cost	e	{10}
Late-delivery cost	l	{1, 10, 50}
Minimum demand	d_{min}	{0}
Maximum demand	d_{max}	{4, 6}
Discount factor	γ	{0.99}

Table 5.2: Averaged optimality gap for all 216 instances

Parameters		Value	Average optimality gap in %	Maximum optimality gap in %
Shipment interval	T	2	0.11	1.79
		4	0.06	1.98
Demand lead time	L_d	1	0.04	1.10
		2	0.13	1.98
Maximum demand	d_{max}	4	0.06	1.47
		6	0.10	1.98
Capacity ratio	r_{Cap}	0.50	0.02	0.43
		1	0.23	1.98
		2	0.00	0.0
Shipment costs	c_2	1	0.03	0.68
		10	0.06	1.47
		50	0.15	1.98
Late-delivery costs	l	1	0.17	1.98
		10	0.08	1.50
		50	0.00	0.04
Total			0.08	1.98

between the DRL policy, π^{DRL}, and the optimal policy computed with VI is presented across varying parameter values in Table 5.2.

Among the 216 instances, π^{DRL} identified the optimal policy in 192 cases. In the few instances where it did not, the difference in optimality was small, as highlighted in Table 5.2. These results highlight the quality of the learning strategy presented, which delivers high performance across a wide range of scenarios.

Table 5.3 lists the maximum number of iterations at which the optimal DRL parameters, denoted by θ^{best}, were identified for each instance. The data reveal that for the majority of instances, θ^{best} was determined in the early phases of the training process. Nevertheless, it is noteworthy that for some instances, updates to θ^{best} occurred toward the end of training, implying that extended training duration may yield marginally better policies.

Table 5.3: Iteration before θ^{best} was found across instances

Iteration	20,000	40,000	60,000	80,000	100,000
Number of instances	170	12	5	12	17

5.4.2 Benchmark Policies

It is highly important to find good approximations to support decisions in large-scale real-world problems. In this section, we provide key observations of the policy obtained by VI, leading to a multi-level threshold policy. The insights offer the development of an ATH policy with three thresholds. Additionally, we present three simple heuristics in this section. The performance of the benchmark policies is tested in Section 5.4.3 for large-scale instances.

Structural Properties of the Observed Threshold Policy

In the following, we describe the structure of the observed threshold policy for a demand lead time L_d. In general, the shipment quantity is composed of four parts:

Part 1 The part of demands that is always shipped. It is limited by the reserved transportation capacity Cap.

Part 2 The part of the remaining demands that is shipped even if the reserved transportation capacity Cap is exceeded.

Part 3 The part of the orders that is shipped. It is limited by the reserved transportation capacity Cap.

Part 4 The part of the remaining orders that is shipped even if the reserved transportation capacity Cap is exceeded.

Let us discuss the intuition leading to these observations. Starting with Part 1, remember that decision-relevant shipment costs for the alternative transportation option, $c_t^{shipment}$, arise only if the inventory manager decides to ship more than Cap units (cf. Equation (5.5)). Additionally, demands cause late-delivery costs c_t^{late} if they are left unfulfilled. Hence, given a specific pre-decision state $\bar{S}_t$, at least $\min(w_t, Cap)$ units are dispatched at the beginning of the shipment period t. If the reserved capacity is exhausted, the *remaining demands* $\bar{w}_t = (w_t - Cap)^+$ may still be shipped at the beginning of the shipment period t. Therefore, a trade-off occurs between minimizing late-delivery and shipment costs, which leads to the observations in Part 2.

The inventory manager has the opportunity to ship orders. Part 3 reflects the case in which there is still available reserved capacity even if the manager decides to ship all the demands w_t. Dispatching orders ahead of their due date causes early-delivery costs, which is why the inventory manager may decide to ship fewer orders than the remaining transportation capacity $(Cap - w_t)^+$. The trade-off occurs between incurring early-delivery costs at the end of the current shipment period versus potentially incurring late-delivery and shipment costs at the end of future shipment periods. Therefore, not necessarily all the orders are shipped until Cap is fully utilized.

Part 4 outlines the case where orders are shipped even if the reserved transportation capacity Cap is fully exhausted. We can compute the number of *remaining orders* that are not shipped with the reserved capacity by $\bar{y}_t^i = \min(y_t^i, (w_t + \sum_{j=1}^{i} y_t^i - Cap)^+) \; \forall \; i \in \{1, 2, \ldots, L_d\}$. In addition to early-delivery costs, shipment costs occur in shipment period t if the manager decides to ship these orders. However, shipping these orders in the current period may still be cost-optimal in the long run as future late-delivery and shipment costs may be saved.

In summary, the structure of the observed threshold policy, denoted by $\pi_{L_d}^{TH}$, can be expressed mathematically as demonstrated in Equation (5.18), including multiple thresholds $\mu_j^i, \forall \; i \in \{1, 2, \ldots, 2L_d\}, \; \forall \; j \in \{0, 1, \ldots, L_d\}$.

$$
\begin{aligned}
\pi_{L_d}^{TH}(\bar{S}_t) = \min\left(w_t, Cap\right) & \\
+ \min\left(\bar{w}_t, \max\left(\left(\left(\bar{w}_t + \sum_{i=1}^{n} y_t^i - \mu_0^n\right)^+\right)_{n \in \{0,1,\ldots,L_d\}}\right)\right) & \\
+ \sum_{i=1}^{L_d}\left(\min\left(\left(Cap - w_t - \sum_{j=1}^{i-1} y_t^j\right)^+, y_t^i, \max\left(\left(\left(\sum_{j=i}^{n} y_t^j - \mu_i^n\right)^+\right)_{n \in \{i,i+1,\ldots,L_d\}}\right)\right)\right) & \\
+ \sum_{i=1}^{L_d}\left(\min\left(\bar{y}_t^i, \max\left(\left(\left(\bar{y}_t^i + \sum_{j=i+1}^{n} y_t^j - \mu_{L_d+i}^n\right)^+\right)_{n \in \{i,i+1,\ldots,L_d\}}\right)\right)\right), & \qquad (5.18)
\end{aligned}
$$

where $\bar{w}_t = (w_t - Cap)^+$ and $\bar{y}_t^i = \bar{y}_t^i = \min(y_t^i, (w_t + \sum_{j=1}^{i} y_t^j - Cap)^+)$ represent the remaining demands and the remaining orders due to i periods, respectively, which cannot be shipped by the reserved capacity Cap.

To emphasize the necessity for entire thresholds, we repeat that orders with distinct due dates cannot be aggregated since they cause different associated costs, leading to $L_d + 1$ state dimensions. The optimal policy obtained by VI demonstrates that the decision of how many units of $\bar{w}_t$ should be shipped at the end of shipment period t depends not only on the number of $\bar{w}_t$ itself but also on the number of orders that may influence this decision. We observe that orders with due dates in the near future have a stronger impact on the abovementioned decision than those with due dates in the distant future. In conclusion, the number of orders cannot be aggregated across several state dimensions, even in context decision-making. Overall, $L_d + 1$ thresholds are needed to represent the influence of the number of orders $y_t^i \; \forall \; i \in \{1, 2, \ldots, L_d\}$ on the dispatch decision for remaining demands $\bar{w}_t$.

The abovementioned impact can be transferred to the dispatch decision for the orders themselves. Specifically, the number of orders $y_t^k \; \forall \; k \in \{1, \ldots, L_d\}$ may have an impact on the shipment decision $y_t^l \; \forall \; l \in \{1, \ldots, k\}$. Consequently, multiple thresholds are needed to represent these influences. Since the influence also depends on whether the shipment incurs additional shipment costs for the alternative transportation option (Part 3 versus Part 4), a total of two $\sum_{k=1}^{L_d} k = \frac{L_d(L_d+1)}{2}$ thresholds must be included. We can conclude that the total number of thresholds in Equation (5.18) grows quadratically according to $(L_d + 1)^2$.

When considering a demand lead time of $L_d = 1$, the total number of thresholds is still tractable at four. The associated policy is shown in Equation (5.19).

$$\begin{aligned}
\pi_1^{TH}(\bar{S}_t) = {} & \min\left(w_t, Cap\right) \\
& + \min\left(\bar{w}_t, \max\left(\left(\bar{w}_t - \mu_0^0\right)^+, \left(\bar{w}_t + y_t^1 - \mu_0^1\right)^+\right)\right) \\
& + \min\left(\left(Cap - w_t\right)^+, y_t^1, \left(y_t^1 - \mu_1^1\right)^+\right) \\
& + \left(\bar{y}_t^1 - \mu_2^1\right)^+
\end{aligned} \tag{5.19}$$

where $\bar{w}_t = (w_t - Cap)^+$ and $\bar{y}_t^1 = \min(y_t^1, (w_t + y_t^1 - Cap)^+)$ represent the remaining demands and the remaining orders, respectively, which cannot be shipped by the reserved capacity Cap.

We computed the expected total period cost for all 108 small-scale instances where $L_d = 1$ by evaluating every possible combination of the four thresholds, where each threshold can take values from 0 to $d_{max}(L_d + 1)$. We then identify the threshold values that yield the minimum expected total period cost and compare it to the results calculated by VI. It is noteworthy that we can determine a combination of thresholds that minimizes the expected total cost per period for all 108 instances.

Approximate Threshold Policy - π^{ATH}

Let us introduce an ATH policy to compute the outbound shipment quantities based on the observed properties of the previous section, and we denote it by π^{ATH}. The main difference from the policy π^{TH} lies in the aggregation of orders with different due dates. Consequently, a total of three thresholds are needed to reflect the impact on shipment decisions, as indicated in Part 2, Part 3 and Part 4, independent of the length of the demand lead time. Specifically, π^{ATH} is defined as follows:

$$\begin{aligned}\pi^{ATH}(\bar{S}_t) = &\min(w_t, Cap)\\ &+ \min\left(\bar{w}_t, \left(\bar{w}_t + \sum_{i=1}^{L_d} y_t^i - \mu_0^{ATH}\right)^+\right)\\ &+ \min\left((Cap - w_t)^+, \left(\sum_{i=1}^{L_d} y_t^i - \mu_1^{ATH}\right)^+\right)\\ &+ \left(\bar{y}_t - \mu_2^{ATH}\right)^+, \end{aligned} \tag{5.20}$$

where $\bar{w}_t = (w_t - Cap)^+$ and $\bar{y}_t = (w_t + \sum_{i=1}^{L_d} y_t^i - Cap)^+$ represent the remaining demands and the remaining orders, respectively, which cannot be shipped by the reserved capacity Cap.

While the decision under π^{ATH} is easier to compute than that under π^{DRL}, the challenge is to determine the cost-minimizing values for the three thresholds μ_0^{ATH}, μ_1^{ATH}, and μ_2^{ATH}. For this purpose, we use a version of simulated annealing, as described in Van Laarhoven et al. (1987), where each threshold can take non-negative values. A detailed description of the simulated annealing algorithm is presented in Appendix 5.6.1.

Simple Benchmark Policies

In addition to the ATH policy, π^{ATH}, we present simple benchmark policies without any thresholds, as detailed in the following:

Dispatch-up-to policy	Under this policy, we dispatch all demands even if the reserved transportation capacity Cap is exceeded and unplanned transportation capacity is needed. However, we dispatch orders only *up to* the reserved capacity Cap and at a maximum of one shipment period ahead. $\pi^{UPTO}(\bar{S}_t) = \max\left(\min(Cap, w_t + \sum_{i=1}^{\min(L_d,T)} y_t^i), w_t\right)$
Lazy policy	Under this policy, we dispatch all demands. This policy is similar to a situation without any ADI. $\pi^{LAZY}(\bar{S}_t) = w_t$
Greedy policy	Under this strategy, we dispatch all the demands and orders. $\pi^{GRDY}(\bar{S}_t) = w_t + \sum_{i=1}^{L_d} y_t^i$

The dispatch-up-to policy is inspired by the outbound shipment policy presented in Section 4.1. In the following, we present the performance of the four heuristics for 270 large-scale instances.

5.4.3 Large-Scale Instances

In this section, we examine the performance of the presented benchmark policies by comparing the associated expected total period cost to the expected total period cost when using DRL. In addition, we present managerial insights into the value of ADI and capacity planning. Like for small-scale instances, we assume that the number of stochastic orders at the end of each period can be represented by a uniform distribution $\mathcal{U}(d_{\text{min}}, d_{\text{max}})$. The configurations for the large-scale instances are presented in Table 5.5. In contrast to small-scale instances, we consider demand lead times from 0 to 8, fix the maximum demand to 10, and investigate two shipment interval levels, 2 and 6. Let us recall that the reserved transportation capacity is $Cap = r_{Cap}T\frac{d_{min}+d_{max}}{2}$. Overall, we consider 270 distinct large-scale instances.

Before we conduct managerial investigations for large-scale instances based on the expected total cost per period obtained by the DRL policy, we first determine how often

Table 5.5: Problem parameters for large instances

Parameters		Values
Shipment interval	T	{2, 6}
Demand lead time	L_d	{0, 2,. . . , 8}
Capacity ratio	r_{Cap}	{0.5, 1, 2}
Shipment cost	c_2	{1, 10, 50}
Early-delivery cost	e	{10}
Late-delivery cost	l	{1, 10, 50}
Minimum demand	d_{min}	{0}
Maximum demand	d_{max}	{10}
Discount factor	γ	{0.99}

the DRL policy achieves the best performance, defined as leading to the minimum total period cost among all mentioned policies. We specify that the DRL policy performs the best among all the mentioned policies if its associated total cost per period is at most 1 % higher than the minimum total cost per period. In our analysis, we find that the DRL policy exhibits the best performance in 252 out of 270 instances. In only 18 instances, either the ATH or the dispatch-up-to policy outperforms the DRL policy. Across all these instances, the average cost increase compared to the minimum cost is 0.18 %, with a maximum increase of 5.23 %. Given this remarkable performance even in large-scale instances, we conclude that the performance of the DRL policy is acceptable as the underlying policy for managerial insights.

Performance of the Benchmark Policies

This subsection primarily focuses on the performance of the benchmark policies. Therefore, we compare the expected total cost per period of all benchmark policies with that of the DRL policy, which we consider to indicate near-optimal total cost values. In Appendix 5.6.2, we present Table 5.8, which shows the average total cost per period for all the policies under fixed input parameters. In Figure 5.3, we specifically concentrate on the expected total period cost across an increased demand lead time L_d. In Table 5.6, we provide a summary of the parameter settings where each of the benchmark policies performs near-optimally.

First, we mention that in instances where no ADI is available, i.e., when $L_d = 0$, all simple benchmark policies (i.e., dispatch-up-to, lazy, and greedy policies) are designed to ensure the shipment of all demands. However, both the DRL and ATH policies slightly outperform the three simple policies. This suggests that it is cost-beneficial, at least for

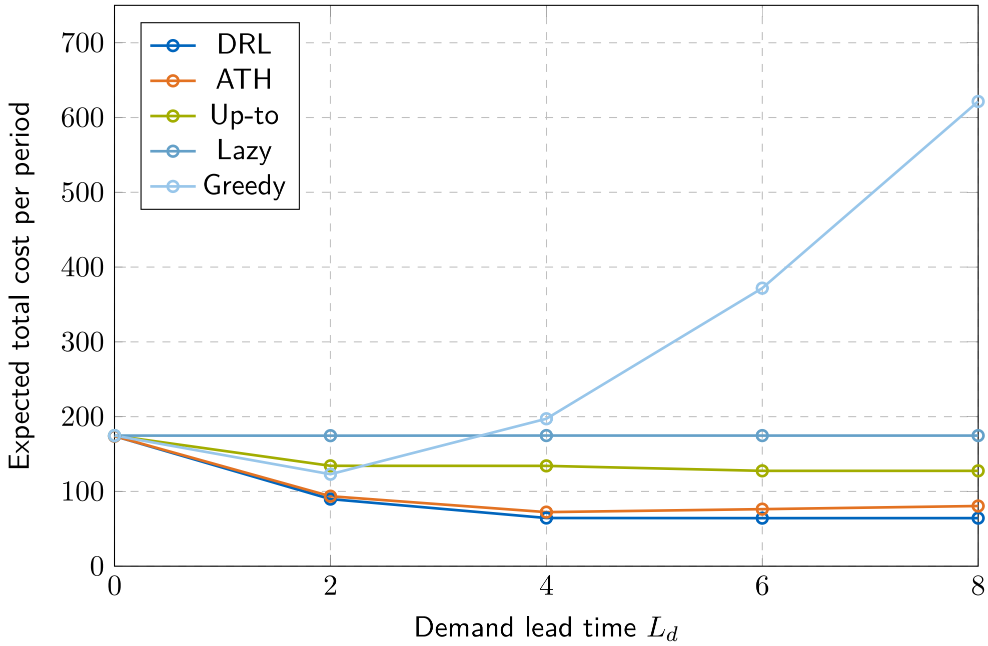

Figure 5.3: Expected total cost per period under different policies across demand lead times

some instances, to delay the fulfillment of demands until the beginning of the subsequent shipment period, as the shipment fees for the alternative transportation option outweigh the costs associated with late deliveries.

The ATH policy closely approximates the performance of the DRL policy and identifies the cost-minimizing policy in more than 65 % of the instances. When no ADI is available, both the DRL and ATH policies minimize the total cost per period. However, the relative total cost difference compared to that of DRL increases as the demand lead time increases. For example, when $L_d = 2$, the three thresholds approximates the near-optimal policy quite well. In 41 out of 54 instances, this approximation finds the near-optimal policy, resulting in an average relative total cost increase of 15.34 % compared to DRL. This high percentage is caused by a few instances leading to a maximum relative cost of 135.87 %.

As the demand lead time increases, e.g., when $L_d = 8$, the relative total cost increase becomes 60.08 %, since 31 out of 54 instances do not perform similarly to the near-optimal policy. In particular, when the shipment interval is set to 2 and the unit late-delivery cost is 50, the three thresholds cannot capture all the influences on near-optimal shipment decisions. However, if we exclude all instances with $T = 2$ and $l = 50$, the average and maximum relative costs increase compared to those of DRL across all remaining instances by only 6.87 % and 25.69 %, respectively. Similarly, we find that the ATH policy works supremely well if the late-delivery cost parameter is 1. To summarize, the ATH policy is an extremely good approximation for most of the instances. An intuitive policy, such as the ATH, is often more interpretable and explainable; therefore, it may be preferable to a DRL policy, which is a black box.

The dispatch-up-to policy already exhibits an average relative total cost increase of more than 100 % across the 54 instances compared to DRL if $L_d = 2$. With increasing demand lead time, the average relative cost increase becomes even greater. However, we find that this policy works extremely well in scenarios where the unit late-delivery cost is high, i.e., $l = 50$, and the capacity ratio is high, i.e., $r_{Cap} = 2$. In these scenarios, a significant number of late deliveries can be avoided by shipping advance orders ahead of schedule, enabled by the large capacity reservation.

The lazy policy acts as a policy without any demand information, which is why the average total cost per period is constant with increasing demand lead time. The lazy policy performs well when the late-delivery cost parameter is low, in combination with a large capacity reservation. Under this policy, only demands that are due are dispatched, which must wait quite long for dispatch due to the time-based shipment scheme. Therefore, this policy seems reasonable only for a low late-delivery cost parameter. Additionally, if a large amount of transportation capacity is reserved, the probability of incurring costs for the alternative transportation option is very low.

The greedy policy ships all the demands and advance orders at the beginning of the shipment period. As early-delivery costs are incurred on a per-unit and per-period basis for units sent in advance, the average total period cost increases with increasing demand lead time under this policy. However, the greedy policy performs extremely well in scenarios where the demand lead time is shorter than the shipment interval and when there is a strong incentive to avoid late deliveries, i.e., when late-delivery costs are high.

In summary, we note that every policy presented in this section has its own justification and is applicable to specific scenarios, summarized in Table 5.6. However, if the planner does not consider the cost parameters or the capacity ratio, then this may lead to poor cost performance.

Table 5.6: Cost minimal behavior of the benchmark policies

	$L_d < T$			$L_d \geq T$		
r_{Cap}	$l = 1$	$l = 10$	$l = 50$	$l = 1$	$l = 10$	$l = 50$
0.5	ATH	-	Greedy	ATH	-	-
1	ATH	-	Greedy	ATH	-	-
2	ATH/Lazy	-	Greedy/Up-to	ATH/Lazy	-	Up-to

Cost Improvements by Advance Demand Information and Flexible Deliveries

ADI facilitates the option to facilitate flexible deliveries and to ship advance orders ahead of schedule to the production facility. In this section, our objective is to emphasize the value of integrating ADI and flexible deliveries by focusing on the average total cost per period under fixed input parameters. We pay particular attention to the impact of increasing demand lead times.

Table 5.7 illustrates the marginal relative cost differences per period as the demand lead time increases when calculating the expected cost using DRL.

If L_d is equal to 0, then no order information is communicated to the inventory manager, thereby restricting decision-making to the dispatch of units that are already due. These costs serve as reference values, and we increase the demand lead time with a step size of two. When we increase L_d from 0 to 2, we observe a significant reduction in the marginal relative cost of 28.38 % across all instances. However, the marginal relative cost reduction decreases as the demand lead time increases. For example, increasing L_d from 2 to 4 results in an average reduction of 8.95 %. Overall, the relative cost difference when increasing the demand lead time from 0 to 8 is 32.55 %.

We observe that a short demand lead time already offers significant cost improvements. This is because late-delivery and shipment costs associated with alternative transportation options can be avoided by shipping orders in advance. These cost savings outweigh the incurred costs for early deliveries. However, as the demand lead time increases, the potential for further savings on shipment costs decreases, as the utilization of the capacity reservation is already high. Additionally, the costs for early deliveries depend on the number of periods shipped in advance, which is why it is reasonable to ship orders several shipment periods in advance only in a specific situation. Consequently, an additional increase in L_d does not necessarily result in notable total cost reductions. In summary, the ADI value does not follow a linear pattern; instead, it decreases as the demand lead time increases.

Table 5.7: Marginal relative cost difference per period enabled by ADI and flexible deliveries using DRL

Parameter	Value	Marginal relative cost difference in % $L_d = 0 \rightarrow L_d = 2$	Marginal relative cost difference in % $L_d = 2 \rightarrow L_d = 4$	Marginal relative cost difference in % $L_d = 4 \rightarrow L_d = 6$	Marginal relative cost in % $L_d = 6 \rightarrow L_d = 8$	Relative cost difference in % $L_d = 0 \rightarrow L_d = 8$
r_{Cap}	0.5	24.71	7.71	−0.12	0.02	28.18
	1	31.08	10.06	0.84	−0.14	36.26
	2	29.35	9.07	−0.01	−0.12	33.21
c_2	1	29.24	9.00	−0.07	0.01	33.11
	10	29.11	8.83	0.28	−0.25	33.04
	50	26.79	9.01	0.51	0.01	31.49
l	1	1.54	0.20	0.28	−0.14	1.84
	10	20.85	0.53	0.14	0.01	21.24
	50	62.75	26.10	0.29	−0.10	74.57
T	2	26.26	0.38	0.09	−0.04	26.53
	6	30.49	17.51	0.39	−0.12	38.57
Total		28.38	8.95	0.24	−0.08	32.55

We emphasize that the unit costs for late deliveries, in particular, have a substantial impact on cost reductions. When late deliveries are cheap, the value of the ADI becomes relatively small, as the cost savings associated with early deliveries are lower than the cost associated with early deliveries. Conversely, when the unit costs for late deliveries are high, the inventory manager benefits from early deliveries by avoiding costly late deliveries.

When comparing shorter and longer shipment intervals, the numerical results demonstrate greater total cost savings in scenarios with longer shipment intervals. In these latter scenarios, the ADI prevents many late deliveries, which are associated with high late-delivery costs. For example, when $T = 2$, shipments dispatch frequently, which is why even without ADI, the late-delivery costs are moderate. Furthermore, we observe that the inventory manager benefits from a demand lead time of only 2 when $T = 2$. Any additional increase in L_d results in negligible cost savings. Based on this observation, we conclude that it is not advisable to expedite orders by more than one shipment period earlier on average. Dispatching orders too early leads to high early-delivery costs, while no further late-delivery cost savings can be achieved.

In summary, having access to information about future demands can result in substantial cost reductions if the inventory manager effectively uses this information by adapting logistic processes. We specifically concentrate on the possibility of flexible deliveries, meaning that production facility agrees on shipments ahead of their due dates.

Cost Comparison under Different Transportation Capacity Reservations

In this subsection, we draw conclusions regarding how an inventory manager should establish transportation capacity based on the relation between shipment costs for primary and alternative transportation options. To enable a comparison of different capacity ratios, we must incorporate the shipment costs that are not relevant to decision-making, i.e., the shipment/reservation costs for the primary transportation option. These costs are excluded and not relevant in previous analyses. Consequently, we add the costs for reserving transportation capacity, $\frac{c_1 Cap}{T}$, to the total cost per period for each instance.

The reservation costs per unit, c_1, must be reasonably related to the costs of the alternative transportation option. To establish this relation, we introduce the shipment cost ratio $r_c = \frac{c_1}{c_2}$. Given that the capacity for the primary transportation option can be offered at a lower cost (i.e., $c_1 < c_2$) since it is set in advance, we restrict r_c to within the range of 0 to 1 $(0 < r_c < 1)$.

We investigate two scenarios and set $r_c \in \{0.50, 0.75\}$. In Figure 5.4 and Figure 5.5, we present the expected total cost per period averaged across all instances while fixing the shipment costs for the alternative transportation option, the shipment cost ratios, and the capacity ratios.

When the shipment costs for both transportation options are low, the manager does not have significant total cost benefits from a specific capacity ratio; the costs associated with early and late deliveries outweigh. However, as the costs of shipments increase, and especially when the primary transportation option is substantially cheaper than is the alternative option, the manager should opt for a capacity reservation close to the mean demand during a shipment cycle. If the ratio between the shipment costs for reservations and unplanned shipments increases, e.g., $r_c = 0.75$, the incentive for the manager to choose $r_{Cap} = 1$ over $r_{Cap} = 0.50$ decreases. Nevertheless, the inventory manager should avoid reserving excessive capacity, as this may not be used and results in unnecessary reservation costs.

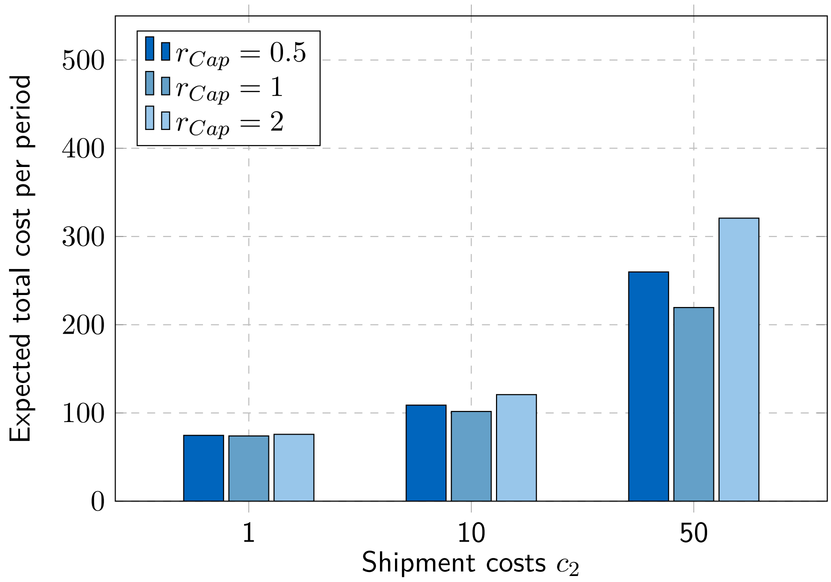

Figure 5.4: Expected total cost across different capacity ratios and shipment cost ratios - $r_c = 0.5$

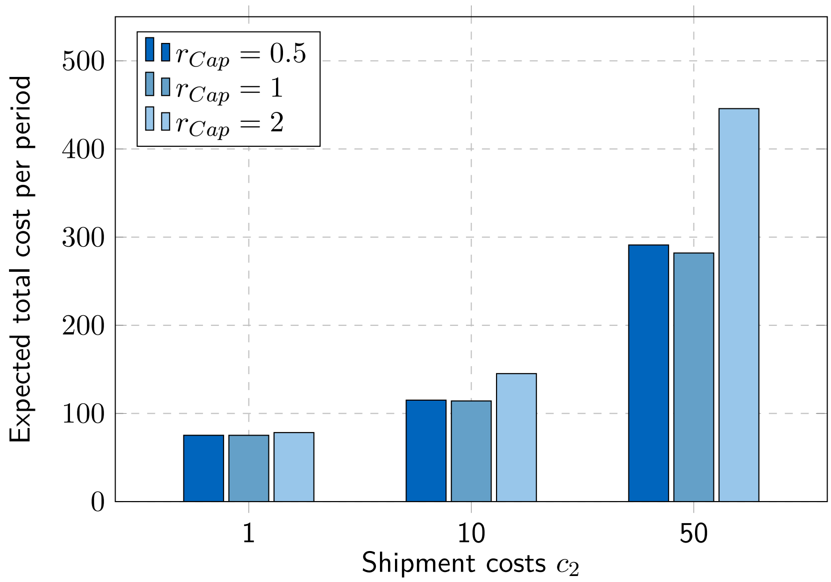

Figure 5.5: Expected total cost across different capacity ratios and shipment cost ratios - $r_c = 0.75$

5.5 Summary and Outlook

In this chapter, we consider a warehouse that supplies a single production facility by using a time-based shipment dispatch scheme. The presence of ADI enables the warehouse manager to dispatch orders ahead of their corresponding due dates. The warehouse manager does not rely on a heuristic shipment policy but instead opts for the optimal shipment policy. We assume that the warehouse is equipped with ample stock, eliminating any potential influence on shipment possibilities due to stockouts. However, we do incorporate late-delivery costs if the inventory manager decides to ship demands after the due date. This reflects the necessity of higher safety stocks at the production facility due to the delayed shipment process. Similarly, early-delivery costs are incurred if orders are satisfied before the due date, representing additional inventory holding costs for the production facility. Finally, we consider shipment costs to distinguish between two distinct transportation modes.

Our numerical study demonstrates that the policy obtained by DRL performs very well for small-scale instances, resulting in an average optimality gap of 0.08 %. Based on the results from value iteration, we observe the structure of the shipment policy, which is a multi-level threshold policy. Given these observations, we present an approximated three-level threshold policy. Additionally, we introduce simple heuristic policies without any thresholds.

For large-scale instances, we conduct a comparative analysis between the DRL shipment policy and various heuristic shipment policies. We find that the ATH shipment policy performs very well for most of the instances. Furthermore, we provide a summary of the input parameter settings under which each policy performs nearly optimally.

In accordance with the previous chapter, We find that the value of ADI does not follow a linear pattern but decreases as the demand lead time increases. Furthermore, in many scenarios, it is advisable for a warehouse manager to reserve transportation capacity within the range of the mean order arrival during a shipment cycle.

We firmly believe that the general structure of the approximate three-level threshold shipment policy yields good performance even in the presence of heterogeneous and imperfect ADI. For future research, we propose applying the ATH policy to a multi-echelon inventory system with time-based dispatching. This involves targeting near-optimal shipment intervals and inventory levels for all stock points while utilizing a heuristic allocation policy such as FCFS.

5.6 Appendix

5.6.1 Implementation of the Simulated Annealing Algorithm

Algorithm 1 shows the pseudocode for the implementation of the simulated annealing algorithm (Van Laarhoven et al., 1987) used to optimize the three thresholds in the ATH approach, namely, μ_1^{ATH}, μ_2^{ATH}, and μ_3^{ATH}.

The algorithm begins by setting the three parameters μ_1^{ATH}, μ_2^{ATH}, and μ_3^{ATH} to zero. The algorithm initializes the temperature to 100 and sets a cooling rate of 0.999, which dictates the probability of exploring suboptimal solutions as the iterations progress. The iteration counters, denoted by $iter$, count the total iterations up to a maximum of $maxIter$ (1,000), and $nonImprovingIter$, count the iterations since the last improvement in the solution.

Two evaluation scores are also calculated at the beginning with the current parameter values: $bestEval$ is the best score found thus far, and $currentEval$ is the score of the current iteration. The main loop of the algorithm iteratively explores new parameter combinations by making small adjustments to the current values, evaluating them, and deciding whether to accept the new combination based on the acceptance probability function.

In this chapter, we perturb μ_1^{ATH}, μ_2^{ATH}, and μ_3^{ATH} (function $neighbor$) by adding a random integer from the set $\{-2, -1, 0, 1, 2\}$ to each threshold, generating $\mu_1^{ATH-new}$, $\mu_2^{ATH-new}$, and $\mu_3^{ATH-new}$. The evaluation of the newly generated parameters (function $eval$) is performed by rolling out the new parameters for 10,000 shipment cycles. The acceptance criterion is determined by the current, the new evaluation scores, and the current temperature (function $accept$). If the new evaluation is better than the current evaluation, then we always accept the newly generated parameters as the current parameter set. Otherwise, we accept it with a probability p, computed as follows:

$$p = \mathrm{e}^{(currentEval - newEval)/temperature} \tag{5.21}$$

When a new combination of parameters yields an improved evaluation score, the algorithm updates the best-known solution ($\mu_1^{ATH-best}$, $\mu_2^{ATH-best}$, and $\mu_3^{ATH-best}$) and resets the counter for non-improving iterations. If the new evaluation does not yield an improvement, the algorithm increases the number of non-improving iterations. To prevent the algorithm

Algorithm 1 Simulated Annealing Algorithm

Initialization:

$\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH} \leftarrow 0, 0, 0$

$temperature \leftarrow 100$

$coolingRate \leftarrow 0.999$

$iter \leftarrow 0$

$nonImprovingIter \leftarrow 0$

$maxIter \leftarrow 1,000$

$bestEval \leftarrow eval(\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH})$

$currentEval \leftarrow eval(\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH})$

Output: $\mu_1^{ATH-best}, \mu_2^{ATH-best}, \mu_3^{ATH-best}$

while $iter < maxIter$ **do**:

  $iter \leftarrow iter + 1$

  $\mu_1^{ATH-new}, \mu_2^{ATH-new}, \mu_3^{ATH-new} \leftarrow neighbor(\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH})$

  $newEval \leftarrow eval(\mu_1^{ATH-new}, \mu_2^{ATH-new}, \mu_3^{ATH-new})$

  $p \leftarrow acceptanceProbability(newEval, currentEval, temperature)$

  if $accept(newEval, currentEval, p)$ **then**

    $\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH} \leftarrow \mu_1^{ATH-new}, \mu_2^{ATH-new}, \mu_3^{ATH-new}$

    $currentEval \leftarrow newEval$

    if $newEval < bestEval$ **then**

      $\mu_1^{ATH-best}, \mu_2^{ATH-best}, \mu_3^{ATH-best} \leftarrow \mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH}$

      $bestEval \leftarrow currentEval$

      $nonImprovingIter \leftarrow 0$

    else

      $nonImprovingIter \leftarrow nonImprovingIter + 1$

    end if

  end if

  if $nonImprovingIter = 20$ **then**

    $nonImprovingIter \leftarrow 0$

    $\mu_1^{ATH}, \mu_2^{ATH}, \mu_3^{ATH} \leftarrow \mu_1^{ATH-best}, \mu_2^{ATH-best}, \mu_3^{ATH-best}$

    $currentEval \leftarrow bestEval$

  end if

  $temperature \leftarrow updateTemperature(temperature, coolingRate)$

end while

from getting stuck in a non-optimal region, if there are 20 consecutive non-improving iterations, the parameters are reset to the best values found thus far.

The temperature, which influences the acceptance probability of new solutions, is updated each iteration according to the cooling rate. This gradual reduction simulates the annealing process, where initially, the algorithm is more likely to explore the solution space (including suboptimal solutions) and becomes more conservative as the temperature decreases, fine-tuning around the best-found solutions. The algorithm halts when it reaches the maximum number of iterations, outputting the best parameters obtained during the process.

5.6.2 Table with the Expected Total Cost per Period

Table 5.8: Expected total cost per period under different policies

Parameters	Value	π^{DRL}	π^{ATH}	π^{UPTO}	π^{LAZY}	π^{GRDY}
r_{Cap}	0.5	122.34	131.42	202.73	204.69	327.68
	1	80.88	86.67	122.01	166.82	289.82
	2	70.78	80.46	88.13	152.37	275.36
c	1	71.86	80.50	118.72	153.46	276.45
	10	81.27	89.69	127.52	163.31	286.30
	50	120.87	128.37	166.63	207.11	330.10
l	1	27.23	27.19	45.98	29.75	257.37
	10	75.30	77.96	88.64	97.20	276.10
	50	171.48	193.40	278.24	396.93	359.38
L_d	0	173.93	173.93	174.66	174.66	174.66
	2	89.73	94.58	134.25	174.56	123.01
	4	64.46	72.31	124.17	174.71	197.17
	6	64.26	76.29	127.52	174.66	371.80
	8	64.29	80.49	127.52	174.56	621.45
T	2	45.47	57.53	61.80	74.70	383.53
	6	137.19	141.51	213.45	274.55	211.70
Average		91.33	99.52	137.62	174.63	297.62

Chapter

6 Summary and Outlook

This thesis aimed to integrate different shipment consolidation policies into inventory management. In the literature, the pure time-based, the pure quantity-based, and the hybrid time-and-quantity-based consolidation strategies have been analyzed in single-echelon inventory systems in great detail. In multi-echelon systems, only pure time-based and quantity-based policies have been examined. Contributing to this field of literature, we analyzed a multi-echelon inventory system with a hybrid consolidation policy.

The idea of combining ADI and consolidation policies in the inventory systems was completely innovative. We have, therefore, developed an inventory model incorporating a heuristic time-based shipment strategy where ADI is used to increase the utilization of the transportation means. Building on this, we analyzed the near-optimal shipment quantities under a time-based policy and ADI. We presented results that illustrate the need for simultaneous inventory and transportation planning. This chapter summarizes the main findings of this thesis by answering the research questions and making suggestions for future research.

6.1 Main Research Insights

Research Question 1: *Under which conditions could a multi-echelon inventory distribution system apply a pure time-based or a pure quantity-based instead of using the dominant hybrid consolidation policy?*

In Chapter 3, we considered an OWMR inventory system that applies a hybrid shipment consolidation policy to replenish the stock of retailer groups facing Poisson demand. Shipments are dispatched according to a time-based schedule; however, additional quantity-based shipments may occasionally be dispatched in between if the number of accumulated demand reaches the consolidation quantity.

The main contribution of this chapter is the derivation to obtain the PMF of the inventory level at each retailer. Based on these PMFs, we can efficiently compute the inventory and shipment costs of the inventory system to evaluate a system and subsequently optimize the inventory and shipment decisions. Our numerical study compares the pure time-based quantity-based policies with the hybrid time-and-quantity-based shipment policy across different instances. By definition, the optimal hybrid policy always outperforms the other two policies. However, when the costs for time-based shipments are significantly lower than for quantity-based shipments, the increase in the expected total cost is moderate when applying a pure time-based instead of the hybrid policy. On the contrary, when the costs for time-based and quantity-based dispatches are equal, a pure quantity-based strategy, i.e., a hybrid policy with shipment interval $T \to \infty$, proves to be optimal.

Research Question 2: *What is the value of incorporating ADI and allowing for flexible deliveries in a single-stage inventory system when satisfying external orders according to a flexible time-based shipment consolidation policy? What is the effect on the optimal inventory and shipment policy parameters?*

In Chapter 4, we derive analytical, approximate total cost expressions to study a single-stage inventory system that receives stochastic orders from a production facility. These orders are satisfied according to a flexible time-based shipment policy. We assume a reserved transportation capacity, which can be extended by a costly transportation option in case of need. The heuristic shipment policy operates as follows: On scheduled time-based shipment days, the stock point has to ship all demands with the primary and the alternative transportation option if sufficient stock on hand is available. Should there be remaining capacity at the primary transportation option, orders will be dispatched until the reserved capacity is fully utilized.

After conducting a comprehensive simulation study to validate the derived mathematical expressions in determining the reorder level and the shipment interval, we found that ADI not only reduces the safety stock needed at the stock point but also enables additional cost savings through the adaption of logistic processes. Our investigation focuses on the modification of the shipment policy by allowing for flexible deliveries. It has been shown that the adaption of the shipment policy leads to the largest part of cost reduction compared to the reduction of safety stocks. Further, we demonstrated that the optimal reorder level at the stock point decreases as the demand lead time increases. The optimal capacity reservation for the primary transportation option, however, remains relatively stable as the demand lead time increases, and should be set around the mean demand during the shipment interval.

Research Question 3: *What is the (near-)optimal structure of the outbound shipment policy under time-based dispatching in inventory systems considering ADI and flexible deliveries?*

In Chapter 5, we introduced a MDP to model an inventory system receiving stochastic customer orders with corresponding due dates, and where shipments are dispatched on a time-based schedule. We considered a primary and an alternative (more expensive) transportation option.

The main objective was to draw conclusions on the structure of the optimal outbound shipment quantities. Our proposed DRL algorithm successfully determined near-optimal decisions for small-scale instances. Based on the near-optimal results for large-scale instances, we observed that the near-optimal shipment quantities can be described by a multi-level threshold policy, with the number of thresholds increasing quadratically with respect to the demand lead time. Therefore, we proposed the ATH policy, which showed great performance across the majority of large-scale instances. This approximation relies on three thresholds, taking into account the due dates and the transportation options used, both of which influence the total cost. Furthermore, we proposed three simple policies to compute the shipment quantities and found that each policy performs very well under specific cost conditions. For example, in scenarios where late deliveries are expensive and should be avoided, a policy should be applied where all demands are dispatched, and additionally orders can be satisfied ahead of their due date if there is remaining capacity at the primary transportation option.

6.2 Further Research Opportunities

There are several promising directions for future research. In Chapter 3 and Chapter 4, the considered demand process is limited to a single unit demand per order. Therefore, an extension could be to incorporate a compound Poisson demand process. However, it is worth noting that our presented approaches already require significant computational effort even with a (pure) Poisson process. For that reason, it could be advisable to develop efficient heuristics that obtain sufficiently good decisions for inventory and shipment policies.

In Chapter 3, we focused on shipment consolidation within an OWMR inventory system. This provides already valuable insights about the dependencies of inventory and shipment decisions, however, there remains potential for analyzing more general multi-echelon inventory systems in the future. However, exact solutions for such systems often prove to be intractable. Consequently, the development of efficient heuristics for OWMR inventory systems are of enormous importance, as they could provide a framework for analyzing general inventory systems.

In Chapter 4 and Chapter 5, we assumed perfect ADI for our models. However, in realistic settings it is likely that customers may cancel their order or modify their order quantities after placement. Therefore, future research could develop inventory models including shipment consolidation and imperfect ADI, for example, by including a probability for order cancellations. Moreover, instead of assuming identical demand lead times for all customers, valuable insights could be obtained by including non-identical demand lead times. However, in this case crossovers may appear why the allocation of orders needs to be taken into account.

The research presented in this thesis does not include fees for receiving ADI from customers. However, customers are especially motivated to share future information when incentives as a bonus or a price reduction are offered. Therefore, a further research direction could be to introduce a cost component for ADI. This would not only increase the authenticity of the model, but also provide practical insights into the optimal amount of information an inventory system should pay for, aiming to minimize the system's total cost.

For decades, inventory systems integrating shipment consolidation or ADI have been analyzed. This thesis contributed to this stream of research by jointly optimizing inventory and shipment decisions across different shipment strategies combined with ADI. Nevertheless, as discussed in the previous paragraphs, open questions remain that require appropriate scientific analysis.

Bibliography

Ahmadi, T., Atan, Z., de Kok, T., and Adan, I. (2019a). Optimal control policies for assemble-to-order systems with commitment lead time. *IISE Transactions* 51(12), 1365–1382.

Ahmadi, T., Atan, Z., de Kok, T., and Adan, I. (2019b). Optimal control policies for an inventory system with commitment lead time. *Naval Research Logistics* 66(3), 193–212.

Ahmadi, T., Atan, Z., de Kok, T., and Adan, I. (2020). Time-based service constraints for invenotry systems with commitment lead time. *OR Spectrum* 42(0), 355–395.

Angelus, A. and Özer, Ö. (2016). Knowledge you can act on: Optimal policies for assembly systems with expediting and advance demand information. *Operations Research* 64(6), 1338–1371.

Axsäter, S. (1990). Simple solution procedures for a class of two-echelon inventory problems. *Operations Research* 38(1), 64–69.

Axsäter, S. (1998). Evaluation of installation stock based (R,Q)-policies for two-level inventory systems with Poisson demand. *Operations Research* 46(3), 135 –145.

Axsäter, S. (2000). Exact analysis of continuous review (R,Q) policies in two-echelon inventory systems with compound Poisson demand. *Operations Research* 48(5), 686–696.

Axsäter, S. (2001). A note on stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science* 47(9), 1306–1310.

Axsäter, S. (2015). *Inventory control*. Springer.

Bookbinder, J. H. and Higginson, J. K. (2002). Probabilistic modeling of freight consolidation by private carriage. *Transportation Research Part E: Logistics and Transportation Review* 38(5), 305 –318.

Bourland, K. E., Powell, S. G., and Pyke, D. F. (1996). Exploiting timely demand information to reduce inventories. *European Journal of Operations Research* 92, 239–253.

Boute, R. N., Gijsbrechts, J., Van Jaarsveld, W., and Vanvuchelen, N. (2022). Deep reinforcement learning for inventory control: A roadmap. *European Journal of Operational Research* 298(2), 401–412.

Çetinkaya, S. and Bookbinder, J. H. (2003). Stochastic models for the dispatch of consolidated shipments. *Transportation Research Part B: Methodological* 37(8), 747 –768.

Çetinkaya, S. and Lee, C.-Y. (2000). Stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science* 46(2), 217–232.

Çetinkaya, S., Tekin, E., and Lee, C.-Y. (2008). A stochastic model for joint inventory and outbound shipment decisions. *IIE Transactions* 40(3), 324–340.

Çetinkaya, S., Mutlu, F., and Wei, B. (2014). On the service performance of alternative shipment consolidation policies. *Operations Research Letters* 42(1), 41 –47.

Çetinkaya, S., Mutlu, F., and Lee, C.-Y. (2006). A comparison of outbound dispatch policies for integrated inventory and transportation decisions. *European Journal of Operational Research* 171(3), 1094–1112.

Chen, F. (2001). Market segmentation, advanced demand information, and supply chain performance. *Manufacturing & Service Operations Management* 3(1), 53–67.

Chen, F. Y., Wang, T., and Xu, T. Z. (2005). Integrated inventory replenishment and temporal shipment consolidation: A comparison of quantity-based and time-based models. *Annals of Operations Research* 135(1), 197–210.

De Moor, B. J., Gijsbrechts, J., and Boute, R. N. (2022). Reward shaping to improve the performance of deep reinforcement learning in perishable inventory management. *European Journal of Operational Research* 301(2), 535–545.

Doherty, S. and Hoyle, S. (2009). Supply chain decarbonization: The role of logistics and transport in reducing supply chain carbon emissions. *World Economic Forum, Geneva*.

Du, B. and Larsen, C. (2017). Reservation policies of advance orders in the presence of multiple demand classes. *European Journal of Operational Research* 256(2), 430 –438.

Figueira, G., van Jaarsveld, W., Amorim, P., and Fransoo, J. C. (2023). The impact of committing to customer orders in online retail. *Manufacturing & Service Operations Management* 25(1), 307–322.

Gallego, G. and Özer, O. (2001). Integrated replenishment decisions with advance demand information. *Management Science* 47(10), 1344–1360.

Gallego, G. and Özer, Ö. (2003). Optimal replenishment policies for multiechelon inventory problems under advance demand information. *Manufacturing & Service Operations Management* 5(2), 157–175.

Gijsbrechts, J., Boute, R. N., Van Mieghem, J. A., and Zhang, D. J. (2022). Can deep reinforcement learning improve inventory management? Performance on lost sales, dual-sourcing, and multi-echelon problems. *Manufacturing & Service Operations Management* 24(3), 1349–1368.

Gradštejn, I. S., Ryžik, I. M., and Jeffrey, A. (1994). *Table of integrals, series, and products*. Academic Press.

Hariharan, R. and Zipkin, P. (1995). Customer-order information, leadtimes, and inventories. *Management Science* 41(10), 1599–1607.

Higginson, J. K. and Bookbinder, J. H. (1994). Policy recommendations for a shipment-consolidation program. *Journal of Business Logistics* 15(1), 87–112.

Higginson, J. K. and Bookbinder, J. H. (1995). Markovian decision processes in shipment consolidation. *Transportation Science* 29(3), 242–255.

Howard, C. (2013). New allocation policies for divergent inventory systems with real-time information and shipment consolidation. *Real-time Allocation Decisions in Multi-echelon Inventory Control, Doctoral Thesis*. Lund University.

Howard, C. and Marklund, J. (2011). Evaluation of stock allocation policies in a divergent inventory system with shipment consolidation. *European Journal of Operational Research* 211(2), 298 –309.

Iglehart, D. L. (1963). Optimality of (s,S) policies in the infinite horizon dynamic inventory problem. *Management science* 9(2), 259–267.

Johansson, L., Sonntag, D. R., Marklund, J., and Kiesmüller, G. P. (2020). Controlling distribution inventory systems with shipment consolidation and compound Poisson demand. *European Journal of Operational Research* 280(1), 90 –101.

Kiesmüller, G. P. and de Kok, A. G. (2005). A multi-item multi-echelon inventory system with quantity-based order consolidation. *BETA publicatie: Working papers, No. 147. Eindhoven: Technische Universiteit Eindhoven.*

Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Lu, Y., Song, J.-S., and Yao, D. D. (2003). Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. *Operations Research* 51(2), 292–308.

Malmberg, F. and Marklund, J. (2023). Evaluation and control of inventory distribution systems with quantity based shipment consolidation. *Naval Research Logistics* 70(2), 205–227.

Malmberg, F., Ralfs, J., Marklund, J., and Kiesmüller, G. P. (2024). Managing inventories in sustainable multi-echelon distribution systems with hybrid shipment consolidation. *Working Paper*.

Marklund, J. (2006). Controlling inventories in divergent supply chains with advance-order information. *Operations Research* 54(5), 988–1010.

Marklund, J. (2011). Inventory control in divergent supply chains with time-based dispatching and shipment consolidation. *Naval Research Logistics* 58(1), 59–71.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518(7540), 529–533.

Mutlu, F., Çetinkaya, S., and Bookbinder, J. H. (2010). An analytical model for computing the optimal time-and-quantity-based policy for consolidated shipments. *IIE Transactions* 42(5), 367–377.

Mutlu, F. and Çetinkaya, S. (2010). An integrated model for stock replenishment and shipment scheduling under common carrier dispatch costs. *Transportation Research Part E: Logistics and Transportation Review* 46(6), 844–854.

OECD (2021). *ITF Transport Outlook 2021*. OECD. https://doi.org/10.1787/16826a30-en.

OECD (2023). *ITF Transport Outlook 2023*. OECD. https://doi.org/10.1787/b6cc9ad5-en.

Oroojlooyjadid, A., Nazari, M., Snyder, L. V., and Takáč, M. (2022). A deep q-network for the beer game: Deep reinforcement learning for inventory optimization. *Manufacturing & Service Operations Management* 24(1), 285–304.

Özer, Ö. (2003). Replenishment strategies for distribution systems under advance demand information. *Management Science* 49(3), 255–272.

Özer, Ö. and Wei, W. (2004). Inventory control with limited capacity and advance demand information. *Operations Research* 52(6), 988–1000.

Papadaki, K. P. and Powell, W. B. (2003). An adaptive dynamic programming algorithm for a stochastic multiproduct batch dispatch problem. *Naval Research Logistics* 50(7), 742–769.

Papier, F. (2016). Supply allocation under sequential advance demand information. *Operations Research* 64(2), 341–361.

Powell, W. B. (2022). *Reinforcement learning and stochastic optimization: A unified framework for sequential decisions*. John Wiley & Sons.

Ralfs, J. and Kiesmüller, G. P. (2022). Inventory management with advance demand information and flexible shipment consolidation. *OR Spectrum* 44(4), 1009–1044.

Ralfs, J., Pham, D. T., and Kiesmüller, G. P. (2024). Optimal outbound shipment policy for an inventory system with advance demand information. *Working Paper*.

Satır, B., Erenay, D. S., and Bookbinder, J. H. (2018). Shipment consolidation with two demand classes: Rationing the dispatch capacity. *European Journal of Operations Research* 270(1), 171–184.

Scarf, H., Arrow, K, Karlin, S, and Suppes, P (1960). The optimality of (S,s) policies in the dynamic inventory problem. *Optimal pricing, inflation, and the cost of price adjustment*, 49–56.

Schoenmeyr, T. and Graves, S. C. (2009). Strategic safety stocks in supply chains with evolving forecasts. *Manufacturing & Service Operations Management* 11(4), 657–673.

Silver, E., Pyke, D., and Thomas, D. (2016). *Inventory and production management in supply chains, Fourth Edition*. Taylor & Francis.

Sonntag, D. R., Schrotenboer, A. H., and Kiesmüller, G. P. (2023). Stochastic inventory routing with time-based shipment consolidation. *European Journal of Operational Research* 306(3), 1186–1201.

Stenius, O., Karaarslan, A. G., Marklund, J., and de Kok, A. G. (2016). Exact analysis of divergent inventory systems with time-based shipment consolidation and compound Poisson demand. *Operations Research* 64(4), 906–921.

Stenius, O., Marklund, J., and Axsäter, S. (2018). Sustainable multi-echelon inventory control with shipment consolidation and volume dependent freight costs. *European Journal of Operational Research* 267(3), 904 –916.

Tan, T., Güllü, R., and Erkip, N. (2007). Modelling imperfect advance demand information and analysis of optimal inventory policies. *European Journal of Operational Research* 177(2), 897–923.

Tan, T., Güllü, R., and Erkip, N. (2009). Using imperfect advance demand information in ordering and rationing decisions. *International Journal of Production Economics* 121(2), 665–677.

Tempelmeier, H. (May 2006). *Horst Tempelmeier, Bestandsmanagement in Supply Chains, 2. Aufl., Norderstedt (Books on Demand) 2006*.

Thonemann, U. (2015). *Operations Management: Konzepte, Methoden und Anwendungen*. Pearson Always learning. Pearson.

Thonemann, U. W. (2002). Improving supply-chain performance by sharing advance demand information. *European Journal of Operational Research* 142(1), 81 –107.

Tijms, H. C. (2003). *A first course in stochastic models*. John Wiley & Sons Ltd.

Toktay, L. B. and Wein, L. M. (2001). Analysis of a forecasting-production-inventory system with stationary demand. *Management Science* 47(9), 1268–1281.

Topan, E., Tan, T., van Houtum, G.-J., and Dekker, R. (2018). Using imperfect advance demand information in lost-sales inventory systems with the option of returning inventory. *IISE Transactions* 50(3), 246–264.

Van Hezewijk, L., Dellaert, N., Woensel, T. V., and Gademann, N. (2023). Using the proximal policy optimisation algorithm for solving the stochastic capacitated lot sizing problem. *International Journal of Production Research* 61(6), 1955–1978.

Van Laarhoven, P. J., Aarts, E. H., van Laarhoven, P. J., and Aarts, E. H. (1987). *Simulated annealing*. Springer.

Vanvuchelen, N., Gijsbrechts, J., and Boute, R. (2020). Use of proximal policy optimization for the joint replenishment problem. *Computers in Industry* 119, 103239.

Veinott Jr, A. F. (1966). On the Opimality of (s,S) Inventory policies: New conditions and a new proof. *SIAM Journal on Applied Mathematics* 14(5), 1067–1083.

Wang, T. and Toktay, B. L. (2008). Inventory management with advance demand information and flexible delivery. *Management Science* 54(4), 716–732.

Wei, B., Çetinkaya, S., and Cline, D. B. H. (2023). Inbound replenishment and outbound dispatch decisions under hybrid shipment consolidation policies: An analytical model and comparison. *Transportation Research Part E: Logistics and Transportation Review* 175, 103135.

Wolff, R. W. (1982). Poisson arrivals see time averages. *Operations Research* 30(2), 223–231.