Accepted Manuscript

Shipment Consolidation with Two Demand Classes: Rationing the Dispatch Capacity

Benhür Satır, Fatih Safa Erenay, James H. Bookbinder

Please cite this article as: Benhür Satır, Fatih Safa Erenay, James H. Bookbinder, Shipment Consolidation with Two Demand Classes: Rationing the Dispatch Capacity, *European Journal of Operational Research* (2018), doi: 10.1016/j.ejor.2018.03.016

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- We analyze how to consolidate two-classes of shipments and ration dispatch capacity

- We minimize shipment & holding costs using a continuous-time Markov decision process

- The optimal policies are of control limit type under particular conditions

- Using these structural properties, we propose an alternative solution approach

- The proposed approach leads to improvements in two real-life cases

1

Shipment Consolidation with Two Demand Classes: Rationing the Dispatch Capacity

Benhür Satır[a], Fatih Safa Erenay[b], James H. Bookbinder[b]

[a]*Department of Industrial Engineering, Çankaya University, Yukarıyurtçu Mah., Mimar Sinan Cad. No:4, Etimesgut, Ankara, Turkey 06970*
[b]*Department of Management Sciences, University of Waterloo, 200 University Avenue West, Waterloo, ON, Canada N2L3G1*

Abstract

We analyze the problem faced by a logistics provider that dispatches shipment orders (parcels or larger packages) of two order classes, viz. expedited and regular. Shipment orders arrive according to a compound Poisson process for each class. Upon an arrival, the logistics provider may continue consolidating arriving orders by paying a holding cost. Alternatively, the provider may dispatch, at a fixed cost, a vehicle containing (a portion of) the load consolidated so far. In addition, the provider must specify the composition of each dispatch by allocating (rationing) the volume of the vehicle between expedited and regular shipment orders. We model this problem as a continuous-time Markov Decision Process and minimize the expected discounted total cost. We prove the existence of quantity-based optimal threshold policies under particular conditions. We also structurally analyze the thresholds of these optimal policies. Based on these structural properties, we develop an efficient solution approach for large problem instances which are difficult to solve using the conventional policy-iteration method. For two real-life applications, we show that the quantity-based threshold policies derived using the proposed approach outperform the time policies used in practice.

Keywords: Logistics, Shipment consolidation, Capacity rationing, Markov decision process, Threshold policies.

1. Introduction

This paper focuses on the problem of allocating (rationing) vehicle capacity between different shipment order types which are dynamically consolidated (to be shipped to-

---

*Corresponding Author: Fatih Safa Erenay; Email: ferenay@uwaterloo.ca; Phone: +1 519 8884567x32521

gether) to save on transportation costs. Transportation is a key part of every supply chain. Cutting costs in the entire chain brings competitive advantage which is vital for both profitability and sustainability. Wilson (2015) reports that transportation costs account for about 5.2% of the nominal gross domestic product of the US in 2014. Therefore, developing methods to decrease transportation costs is desirable and may have significant impact on a national economy. For a courier company or a third-party logistics provider (3PL), reducing transportation costs directly results in enhanced margins, i.e., more profit per shipment.

Courier companies and 3PLs offer several service options with distinct prices. For example, let us consider a one-ounce letter to be sent from New York City to Miami. Among some of the choices that UPS offers are **UPS Next Day Air** with guaranteed delivery in one day, one of the fastest shipping services, for $75.48, and **UPS Ground**, an ideal service with delivery guaranteed in five days for a price of $9.90[1]. The price ratio of expedited to normal services in this example is 7.62; however, these ratios may vary, based on proximity of the shipment zones and the transportation medium. For example, UPS Turkey's price ratio of **emergency express** to **standard delivery** on the İstanbul-Ankara route is around 5. On the other hand, this price ratio is 12.7 for MNG Kargo, which frequently ships parcels between Ankara and İstanbul via plane[2].

There are also differences in the business operations of the logistics companies. For instance, UPS Turkey uses both their own vehicles and for-hire trucks between the Mahmutbey Hub (in İstanbul) and the hub in Ankara, depending on the demand. Parcels collected at various UPS outlets in İstanbul and Ankara are transferred through these hubs. UPS Turkey applies mainly a time-based vehicle-dispatch policy according to the tentative vehicle schedules. On the other hand, each day MNG Kargo sends between two and ten fixed-weight shipments of particular types of parcels from Ankara to İstanbul by paying a fixed cost for each shipment[3]. Therefore, MNG Kargo needs to dynamically determine the number of dispatches on the Ankara-İstanbul route based on demand realization, as well as the timing and composition. Similar applications exist for inbound logistics. EKOL Logistics, a 3PL serving Turkish manufacturers on defined milk-run routes, develops a shipment plan specifying when and how much to collect from suppliers, according to the production needs of the manufacturers and the degree of urgency.

---

[1]The quoted delivery prices are retrieved from https://wwwapps.ups.com/ctc on June 17, 2016.

[2]Information on MNG Kargo's delivery prices were retrieved from http://service.mngkargo.com.tr/musteri/WebFaturaListesi/UcretHesapla2.aspx on Nov.15, 2015.

[3]Information on MNG Kargo's system was gathered by phone interview with the Customer Relations and Tele-Marketing Manager of MNG Kargo on July 23, 2012.

Although the preceding logistics companies employ different business models, each faces a similar challenge. They could dispatch vehicles frequently to lower the holding costs and improve customer satisfaction. Alternatively, they may further consolidate the arriving shipment orders, to increase the vehicle utilization and obtain economies of scale on the fixed dispatching-cost. In addition, when dispatching a vehicle, they also need to decide the composition of the load in terms of order types. Because most logistics companies can continuously monitor their orders, those decisions can be made dynamically to improve the system performance.

Shipment consolidation aims to increase vehicle utilization by combining two or more shipment orders, dispatched as an aggregate unit. Recent surveys imply that the majority of American manufacturers use shipment consolidation as an outsourced logistics function to cut costs, and most large 3PLs provide freight consolidation services (e.g. Lieb and Lieb (2015)). These results show that shipment consolidation is a powerful tool for logistics providers. Ülkü (2012) points out that shipment consolidation also helps in achieving "green" supply chain targets by reducing energy waste and carbon emission. When the decision is made to dispatch a vehicle to carry (a portion of) the consolidated load, the next thing to decide is "*How much of each shipment order type should be put on the vehicle?*". Allocating the available volume of the transportation medium to different order types is a form of "**capacity rationing**".

In this context, we propose a continuous-time Markov Decision Process (MDP) model to optimize the decisions on consolidation and vehicle capacity rationing upon the arrival of each shipment order. We consider the perspective of the logistics provider (L-P), whose objective is to minimize total expected discounted cost: sum of transportation plus holding cost. We assume that shipment orders arrive according to a compound Poisson process, and orders are either for **expedited (Type 1)** or **regular (Type 2)** shipments. Although consideration of only two classes of orders is limiting, it is a valid assumption for particular business models. For instance, courier providers in Turkey usually offer options of delivery in one day or in two days, on the main routes within the country.

To the best of our knowledge, this is the first work to jointly analyze the decisions on dispatch timing and load composition from the perspective of shipment consolidation and vehicle capacity rationing. Benefiting from and building upon existing publications, we structurally analyze the optimal solutions of this problem. We show the existence of optimal quantity-based threshold policies for particular cases and further characterize those optimal solutions. Moreover, we utilize the proven structural properties to develop efficient solution approaches for large problem instances. Next, the model is applied to

two real-life cases: those of UPS Turkey and EKOL Logistics. Our numerical experiments show that the quantity-based threshold policies outperform the currently used time policies.

2. Literature Review

Our work is related to the literature on shipment consolidation, customer rationing, and multiproduct batch-service problems. The studies on shipment consolidation mainly analyze for how long (i.e., time policy), or up to what quantity or weight (i.e., quantity-based policy), the shipment orders should be accumulated before a consolidated load is dispatched. Most of this literature aims to minimize the total transportation and holding cost, assuming Poisson-distributed shipment order arrivals.

In what follows, it will be important to distinguish between the cases of "common carriage" and "private carriage". A common carrier is a public, for-hire transportation provider (e.g., trucking company). Private carriage refers to transportation in one's own vehicle, i.e. a truck owned or controlled by the shipper of the goods. Higginson and Bookbinder (1994) employ discrete-event simulation to evaluate performances of particular time, quantity, and time-and-quantity policies for shipment consolidation problems with common carriage under different parameter settings. Focusing on quantity-based policies, Higginson and Bookbinder (1995) propose an MDP model for the consolidation of random-size shipment orders. They illustrate that, for the private-carriage setting, the optimal policy is of control-limit type, which may not be the case for the common-carriage setting. Bookbinder and Higginson (2002) also employ a stochastic renewal-process model for this problem, to derive effective time-and-quantity (hybrid) policies for transportation by private carriage.

A series of papers study the shipment consolidation problem with more general order-arrival processes or alternative performance metrics. Using renewal theory, Çetinkaya and Bookbinder (2003) derive explicit expressions for the optimal quantity-based and time policies for a shipment consolidation problem under private carriage for a general arrival process. They also propose approximate methods to derive effective time-and-quantity policies for the case of common carriage. Bookbinder et al. (2011) propose a Markov-process model with a Markovian phase-type batch-arrival process for a private carrier with infinite dispatch capacity. They develop efficient algorithms for computing the performance measures of quantity, time, and hybrid policies. Merrick and Book-binder (2010) evaluate performances of quantity, time, and time-and-quantity policies, considering both environmental impacts and profitability. Çetinkaya et al. (2014) derive

analytical results comparing the performances of shipment consolidation policies using service-based criteria such as maximum waiting time and average order delay.

There are also a few studies that analyze shipment consolidation as a process mechanism in larger problems on shipment planning for airfreight forwarders (Wong et al. 2009) and 3PLs (DallOrto et al. 2006, Ülkü and Bookbinder 2012), as well as material routing in inventory networks (Howard and Marklund 2011).

The present paper is also related to the literature on rationing. Those publications analyze how the available inventory or service capacity should be rationed between different classes of customers to maximize total profit or benefit. For instance, Ha (1997) considers the problem of rationing finished-goods inventory for a manufacturer of a single item in a make-to-stock system, with several demand classes (different selling prices) and lost sales, under the assumption of exponential manufacturing and demand interarrival times. He derives the optimality equations for solving this problem and proves the existence of optimal threshold-type policies. de Véricourt et al. (2000) study the problem of rationing the production capacity of a manufacturer between two products and characterize the optimal solution under particular conditions. Yang and Fung (2014) model a manufacturer in a make-to-stock environment facing uncertainty on both the demand and supply sides, using a finite-horizon MDP. They show characteristics of the optimal admission policy for multiple customer classes. In the vehicle-rental industry, Pazour and Roy (2015) consider a system serving both priority and non-priority customers using a pool of homogeneous vehicles. Focusing on only threshold-type policies, they employ a queueing model to obtain exact solutions for the best threshold selection. Further references on inventory rationing are contained in Fadıloğlu and Bulut (2010).

Finally, our work is related to batch-service problems in which a decision maker determines when to serve accumulating orders or customers together as a group. Papadaki and Powell (2002) propose a finite-horizon MDP model for such problems, where homogeneous customers arrive in batches of random size according to a Poisson process. If the decision is to serve, customers are processed up to a particular service capacity within the current time epoch. Their MDP model minimizes the total holding and fixed service costs. Papadaki and Powell (2002) show that the minimum value function is monotone non-decreasing, and that the optimal policy is of the control-limit type. Papadaki and Powell (2003, 2007) extend these results to a non-homogeneous customer setting (e.g., n types of customers), and prove that the optimal policy is either wait or serve, by sequencing the customers based on their holding costs. Min (2014) also proposes an infinite-horizon MDP model for a multi-class batch service problem with class-dependent waiting costs. This model considers a time- and batch-size-dependent

variable service cost incurred when the total service time of the current batch exceeds the server's shift time (e.g., overtime cost). Min (2014) analyses the structure of the optimal solution and proposes heuristic approaches to solve the problem.

Our research differs from publications in the shipment consolidation literature because they do not consider multiple classes of shipment orders (e.g., expedited vs. regular shipments). Although studies in the rationing literature consider multiple order/demand classes, those analyses process the orders as they arrive rather then processing them as a batch. Papadaki and Powell (2003) and Papadaki and Powell (2007) are the works most relevant to ours. We show that our proposed continuous-time discounted-cost MDP model can be simplified to an infinite-horizon discounted-cost MDP. This simplified model is actually equivalent to the infinite-horizon version of the models in Papadaki and Powell (2003) and Papadaki and Powell (2007). However, our work still contributes to the literature as we extend the modeling framework to an infinite-horizon setting, and further characterize the optimal solutions by proving the existence of threshold-type optimal policies and the monotonicity (non-increasing) of those optimal thresholds in particular cases. These analytical results are important because infinite-horizon MDPs suffer from the curse of dimensionality even if they consider only two actions in each state (e.g., as in the optimal stopping time problems in Alagöz et al. (2004, 2007)). Using the existence of monotone optimal thresholds, we develop alternative solution algorithms, enabling us to solve large problem instances which are hard to solve using the conventional policy-iteration algorithm. Finally, although Papadaki and Powell (2003, 2007) reported that their framework is applicable to logistics problems, ours is the first work to employ such an MDP framework in the setting of shipment consolidation and capacity rationing, using real-life data and examples.

3. Methodology

3.1. The Model

We model this problem as a continuous-time discounted-reward Markov Decision Process (CTMDP) defined over an infinite-horizon. The proposed model reflects the perspective of an L-P offering two types of shipment services. That is, arriving shipment orders require either expedited (expedited orders) or standard shipment (regular orders). We refer to these orders as Type 1 and Type 2, and assume that they arrive according to a compound Poisson process with rates λ_1 and λ_2, respectively. We denote the total arrival rate as $\lambda = \lambda_1 + \lambda_2$. We assume that the size of each shipment order is a discrete random variable where $d_1(n)(d_2(m)), n \in \{1, 2, ..., N\}$ ($m \in \{1, 2, ..., M\}$) denotes the

7

probability that an arriving expedited (regular) order requires shipment of a load of size n (m) units.

In the model, decisions are triggered by arrivals of shipment orders, i.e., the time between consecutive decision epochs is exponentially distributed with rate λ. At each order arrival, the L-P either i) continues to consolidate arriving orders (*WAIT*) or ii) dispatches all or a portion of the consolidated load with a proper transportation vehicle (truck), by deciding how much of each order type to include in the shipment (*SHIP*). We assume that orders of the same type are processed according to the first-come-first-serve rule, after deciding to *SHIP*.

The objective function includes the expected transportation and holding costs. The fixed cost to dispatch or release a single vehicle is denoted by K. The latter category merits further discussion. One drawback of dynamic consolidation is a stochastic time of delivery, because the time until the next shipment depends on the randomly-accumulated shipment load. Therefore, consolidation decisions should consider appropriate lateness measures. For this purpose, we define c_1 and c_2 for Type 1 and Type 2 customers, respectively, as holding costs per unit order per unit time. We assume $c_1 > c_2$ and denote $C = (c_1, c_2)$ as the holding cost vector. These holding costs should not be interpreted in the sense of "ownership" of the transported goods. Rather, the holding costs are proxies for the disutility experienced by the customers whose orders are still waiting to be delivered, as well as for the efforts required to store and maintain these goods. Holding costs implicitly reflect due dates of customer orders in other studies in the literature, e.g., Yılmaz and Savaşaneril (2012).

In our model, the system state at time t, $S_t \in \mathbf{S} \equiv \{(s_1, s_2);\ s_1,\ s_2 \in \mathbb{Z}_{\geq 0}\}$, tracks the amounts (in units) of consolidated expedited (s_1) and regular (s_2) orders awaiting shipment. In state $S_t = S$, $A_t(S) \in \mathbf{A}(S) \equiv \{(a_1, a_2);\ a_1, a_2 \in \mathbb{Z}_{\geq 0},\ a_1 \leq s_1,\ a_2 \leq s_2,\ a_1 + a_2 \leq \omega\}$ represents any feasible action. In this notation, a_1 and a_2 refer to the numbers of expedited and regular orders to be dispatched via a vehicle with enough capacity to carry ω units. Naturally, $a_1 = a_2 = 0$ refers to *WAIT*; whereas, $a_i > 0$ (for any i) refers to the *SHIP* decision. Because interarival times between two consecutive shipment orders are exponentially distributed, the distribution of future events beyond time t after observing $S_t = S$ is equivalent to those beyond time $t + k$ after observing $S_{t+k} = S$. Therefore, the optimal actions for the same system state at any two time points are the same in our infinite-horizon continous-time model. Thus, it is sufficient to consider only stationary decisions, i.e., $A_t(S) = A(S)\ \forall S \in \mathbf{S}$, $t \geq 0$.

When calculating the total cost, we need to keep track of the system state only at time points of state change. Therefore, the decision epochs denote time points at which

8

shipment orders arrive, i.e., epoch $p \in \{0, 1, 2, ...\}$ refers to the p^{th} order which arrived at a random time T_p where $T_0 = 0$. Figure 1 shows the sequence of events occurring at the beginning of decision epoch p, given that there was no shipment in epoch $p-1$ and a shipment is dispatched in the current decision epoch (i.e., $A(S_{T_{p-1}}) = \overrightarrow{0}, A(S_{T_p}) > \overrightarrow{0}$)[4]. In the figure, $S_{T_p} = (s_1, s_2)$ denotes the state at the beginning of epoch p immediately after the new shipment order arrival. Without loss of generality, we assume that L-P observes the size of the new shipment order and the system state S_{T_p}, and makes a shipment dispatch decision $A(S_{T_p}) = (a_1, a_2)$ which takes effect instantaneously. Therefore, the holding cost between epochs p and $p + 1$ is incurred for having $s_1 - a_1$ expedited orders and $s_2 - a_2$ regular orders during the interval (T_p, T_{p+1}).
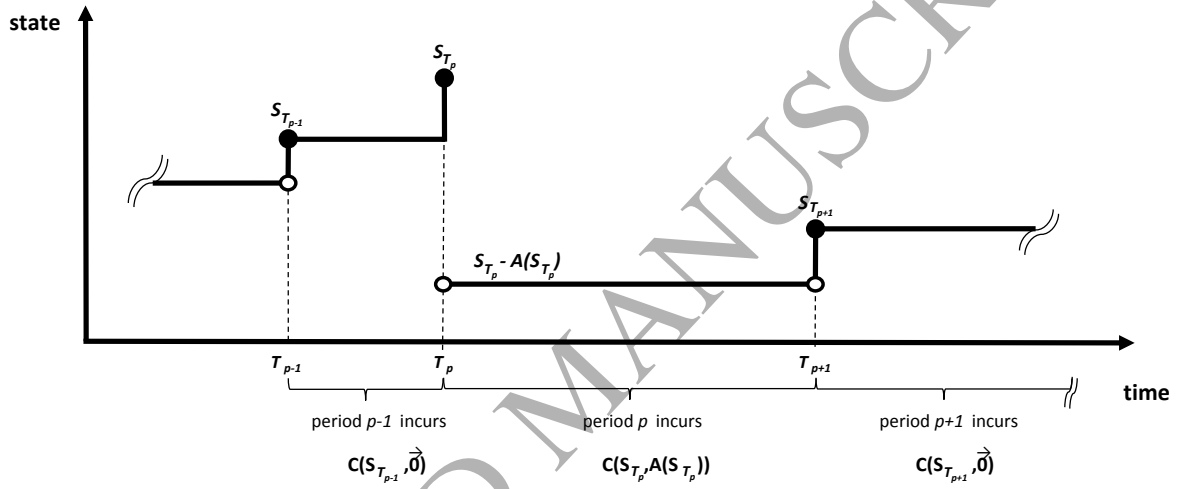

Figure 1: State Evolution over Time

Equation (1) presents the transition probabilities of moving from state $S_{T_p} = S = (s_1, s_2)$ to state $S_{T_{p+1}} = S' = (s_1', s_2')$ under action $A(S_{T_p}) = A = (a_1, a_2) \in \mathbf{A}(S)$ just after an order arrival. Note that λ_1/λ is the probability that the order which arrives in state S is an expedited order. In this case, the state after shipment increases by ne_1 with probability $d_1(n)$. Similar remarks hold for the remaining transition probabilities.

$$
\begin{aligned}
p(S_{T_{p+1}}|S_{T_p}, A(S_{T_p})) &= p(S'|S, A) \\
&= \begin{cases}
d_1(n)\lambda_1/\lambda, & \text{if } (s_1', s_2') = (s_1 - a_1 + n, s_2 - a_2) \quad \forall n \in \{1, ..., N\} \\
d_2(m)\lambda_2/\lambda, & \text{if } (s_1', s_2') = (s_1 - a_1, s_2 - a_2 + m) \; \forall m \in \{1, ..., M\} \\
0, & otherwise
\end{cases}
\end{aligned}
\tag{1}
$$

---

[4] $\overrightarrow{0} = (0, 0)$, $e_1 = (1, 0)$, $e_2 = (0, 1)$

9

Since the objective function is the minimization of discounted total cost, $C(S_{T_p}, A(S_{T_p}))$ represents the immediate cost between two consecutive order arrivals $\forall S_{T_p} = S \in \mathbf{S}$ and $\forall A(S_{T_p}) = A \in \mathbf{A}(S)$ as follows:

$$
\begin{aligned}
C(S_{T_p}, A(S_{T_p})) = C(S, A) &= I_{[A]}K + E\left[\int_0^T [c_1(s_1 - a_1) + c_2(s_2 - a_2)]e^{-\alpha t}dt\right] \quad (2) \\
&= I_{[A]}K + [c_1(s_1 - a_1) + c_2(s_2 - a_2)]/[\alpha + \lambda] \\
&= I_{[A]}K + C(S - A)^{tr}/[\alpha + \lambda]
\end{aligned}
$$

where $\alpha \in (0, 1)$ is the continuous discount rate, $I_{[A]}$ is an indicator function which is equal to 1 if $a_1 + a_2 > 0$, and $(S - A)^{tr}$ is the transpose of the vector representing the state just after the current shipment. Note that the expectation in Equation (2) is w.r.t. the exponentially distributed interarrival time, $T = T_{p+1} - T_p$.

Let π denote any stationary policy, and $A^\pi(S)$ refer to the action for state S under policy π. We define the optimal value function $V(S)$, which represents the minimum expected total discounted cost starting from state S just after an order arrival (i.e., $T_0 = 0$ and $S_0 = S$), in Equations (3) and (4). In these optimality equations, $V(S, A)$ is the expected discounted cost given that action A is chosen at the initial state and the optimal policy is followed from there on, and $\beta = \lambda/(\alpha + \lambda)$.

$$
\begin{aligned}
V(S) &= min_\pi\left\{E\left[\sum_{p=0}^\infty e^{-\alpha T_P}C(S_{T_p}, A^\pi(S_{T_p}))|S\right]\right\} \quad (3) \\
V(S) &= min_{A \in \mathbf{A}(S)}\left\{V(S, A)\right\} \\
&= min_{A \in \mathbf{A}(S)}\left\{I_{[A]}K + \frac{1}{\alpha + \lambda}C(S - A)^{tr} + \beta\sum_{S' \in \mathbf{S}}p(S'|S, A)V(S')\right\} \quad (4)
\end{aligned}
$$

Equation (3) presents the optimality equation for the CTMDP which minimizes the total expected discounted cost accrued through the decision horizon. In Appendix A, we show that Equation (3) can be simplified to Equation (4) following the uniformization procedure described in Equation (11.5.6) of Lippman (1975). Actually, Equation (4) is nothing but the Bellman optimality equation of an embedded discrete-time Markov Decision Process (DTMDP) equivalent to the original CTMDP model. The simplified DTMDP model in Equation (4) is also equivalant to the infinite-horizon version of a special case of the MDP model in Papadaki and Powell (2003, 2007), where the number of product types is equal to two. Equation (4) is defined for any vehicle capacity. In this paper, we will consider two specific capacity cases: **Capacitated Model** with $\omega < \infty$ and **Uncapacitated Model** with $\omega = \infty$.

10

3.2. Structural Analysis

Papadaki and Powell (2003, 2007) show that, for the finite-horizon version of the proposed model, the minimum value function in Equation (4) is monotone and the optimal policy is either to *WAIT* or to *SHIP* the consolidated load by prioritizing the expedited orders. We use these properties to prove that quantity-based optimal threshold policies exist for the uncapacitated model ($\omega = \infty$), and that these thresholds are of "linear staircase" form. All proofs are available in Appendix B. The following definitions introduce the partial-ordering operator and monotonicity type used in this section.

Definition 1. *We define the partial ordering operator $\succeq$ on the two-dimensional set $\Psi = \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ such that $X' \succeq X$ for $X = (x_1, x_2), X' = (x_1', x_2') \in \Psi$, if $x_1' \geq x_1$ and $x_1' + x_2' \geq x_1 + x_2$.*

Definition 2. *A real-valued function F defined on the two-dimensional set Ψ is partially non-decreasing w.r.t. the ordering defined in Definition 1 if we have $F(X') \geq F(X)$ for all $X, X' \in \Psi$ when $X' \succeq X$.*

Theorem 1 states that the optimal value function in Equation (4) is monotone w.r.t the partial ordering defined above. Theorem 1 follows the structural properties presented in Papadaki and Powell (2003, 2007) proving that the optimal value function is monotone non-decreasing in state and shipment order type for the finite-horizon version of our model. These properties from Papadaki and Powell (2003, 2007) apply to the infinite horizon, based on a result in Bertsekas (2001) (page 8). Collectively, these properties imply that $V(S) \leq V(S + k_1 e_1 + k_2 e_2)$ where $S = (s_1, s_2)$, $k_1 \geq 0$, and $k_2 \geq -k_1$, which guarantees the monotonicity property in Theorem 1.

Theorem 1. *$V(S)$ is partially non-decreasing w.r.t. the ordering defined in Definition 1.*

The following lemma illustrates a dominance rule between particular actions. This rule indicates that, when there is room in the vehicle, shipping more and/or replacing regular orders with expedited orders in the shipment load saves cost.

Lemma 1. *Let $A, A' \neq \overrightarrow{0}$ be any pair of feasible actions for state S where $A \succeq A'$, i.e., $A' = A - k_1 e_1 + k_2 e_2$ where k_1 and k_2 are integers such that $a_1 \geq k_1 \geq 0$ and $\min(k_1, s_2 - a_2) \geq k_2 \geq -a_2$. Then, $V(S, A) \leq V(S, A')$.*

Based on Lemma 1, we define the best load composition for a *SHIP* decision, namely action $\overline{A}(S)$, for any $S \in \mathbf{S}$. Action $\overline{A}(S)$ prioritizes expedited orders in utilizing vehicle

11

capacity, i.e., $\overline{A}(S) = (\overline{a}_1, \overline{a}_2)$ such that $\overline{a}_1 = min\{\omega, s_1\}$ and $\overline{a}_2 = min\{\omega - \overline{a}_1, s_2\}$. According to this definition,

$$\overline{a}_1 + \overline{a}_2 = \left\{ \begin{array}{ll} s_1 + s_2, & \text{if } s_1 + s_2 \leq \omega \\ \omega, & \text{if } s_1 + s_2 > \omega. \end{array} \right.$$

Proposition 1 states that the optimal decision for any state S is either to $WAIT$ or to $SHIP$ according to action $\overline{A}(S)$. This proposition is equivalent to a special case of Proposition 4.1 of Papadaki and Powell (2007). Proposition 1 can also be proven by contradiction. Let $A \in \mathbf{A}(S)$ be any feasible action not equal to $\overrightarrow{0}$ or $\overline{A}(S)$. The action $\overline{A}(S)$ ships as many Type 1 orders as vehicle capacity allows and utilizes the vehicle capacity as much as possible. Thus, action A ships either fewer Type 1 orders or dispatches a smaller total load than $\overline{A}(S)$ does. Therefore, the optimality of such an action, i.e., $V(S, A) = V(S) \leq min\{V(S, \overrightarrow{0}), V(S, \overline{A})\}$, contradicts Lemma 1.

Proposition 1. *For any state $S \in \mathbf{S}$, the optimal action can be defined as $A^*(S) = argmin_{A \in \{\overrightarrow{0}, \overline{A}(S)\}}\{V(S, A)\}$, where $\overrightarrow{0}$ refers to WAIT and $\overline{A}(S)$ refers to the SHIP decision, by prioritizing expedited orders while shipping the consolidated load, i.e., $\overline{A}(S) = (min\{s_1, \omega\}, min\{\omega - \overline{a}_1, s_2\})$.*

The minimum value function $V(S)$ in Equation (4) can be simplified as follows, by considering only the two actions specified in Proposition 1 for each state.

$$\begin{aligned} V(S) &= min\Big\{ \frac{1}{\alpha + \lambda}CS^{tr} + \sum_{n=1}^{N}\frac{d_1(n)\lambda_1}{\alpha + \lambda}V(S + ne_1) + \sum_{m=1}^{M}\frac{d_2(m)\lambda_2}{\alpha + \lambda}V(S + me_2), \\ &\quad K + \sum_{n=1}^{N}\frac{d_1(n)\lambda_1}{\alpha + \lambda}V(S' + ne_1) + \sum_{m=1}^{M}\frac{d_2(m)\lambda_2}{\alpha + \lambda}V(S' + me_2)\Big\} \end{aligned} \quad (5)$$

where $S' = (s'_1, s'_2) = \Big(s_1 - min(\omega, s_1), s_2 - min(\omega - min(\omega, s_1), s_2)\Big)$. Proposition 1 is important as it reduces the computational burden for solving the shipment consolidation & capacity rationing problem. However, solving the infinite-horizon DTMDP in Equation (5) is still computationally challenging in real-life instances, due to the curse of dimensionality. To solve this problem efficiently, it is desirable to further characterize the optimal policies. We do so for the uncapacitated model, i.e., $\omega = \infty$, by proving the existence of the optimal threshold-type policies with linear-staircase thresholds. Since $\overline{A}(S) = S$ when $\omega = \infty$, $V(S)$ can be expressed as follows in this case:

12

$$V(S) = min\Big\{\frac{1}{\alpha+\lambda}CS^{tr} + \frac{\lambda_1}{\alpha+\lambda}\sum_{n=1}^{N}d_1(n)V(S+ne_1) + \frac{\lambda_2}{\alpha+\lambda}\sum_{m=1}^{M}d_2(m)V(S+me_2),$$

$$K + \frac{\lambda_1}{\alpha+\lambda}\sum_{n=1}^{N}d_1(n)V(ne_1) + \frac{\lambda_2}{\alpha+\lambda}\sum_{m=1}^{M}d_2(m)V(me_2)\Big\} \tag{6}$$

Theorem 2 states the existence of the optimal infinite-horizon threshold policies when $\omega = \infty$. Our numerical experiments indicate that Theorem 2 may also hold for the capacitated problem instances (i.e., $\omega \leq \infty$). We realize that the proof of Theorem 2 can be extended to the case of $\omega \leq \infty$ for S and S' state couples that are either sufficiently small or sufficiently large. However, we leave a complete generalization of Theorem 2 for future studies.

Theorem 2. *If $\omega = \infty$, then the optimal policy is of control-limit type. That is, if $A^*(S) = S$ for state S, then $A^*(S') = S'$ for any state $S' \geq S$.*

We observe that the structure of the optimal policies can be even nicer in some cases, as shown in Lemma 2 and Theorem 3.

Lemma 2. *Suppose that $\omega = \infty$, and r and q are minimum positive integers satisfying $rc_1 = qc_2$. Then $V(S) = V(S + re_1 - qe_2)$ for any $S = (s_1, s_2)$ s.t. $s_2 \geq q$.*

Based on Lemma 2, Theorem 3 proves the existence of the optimal linear-staircase threshold policy for $\omega = \infty$ when $\frac{q}{r}$ is integer (i.e., $r = 1$). Note that Theorem 3 can be extended to prove the existence of the optimal non-increasing staircase thresholds (though not perfectly linear) when $\frac{q}{r}$ is not an integer, because Lemma 2 is valid for all integer r and q values. This proof is omitted for brevity.

Theorem 3. *Let $\omega = \infty$ and $c_1 = qc_2$ where q is a positive integer. There exists an optimal linear-staircase threshold of the state variables, beyond which the optimal action is to SHIP and below which it is to WAIT. That is, there exists an $\overline{s_2}(s_1) \ \forall s_1 \in \{0, 1, ...\}$ such that i) $\overline{s_2}(s_1 + 1) = \overline{s_2}(s_1) - q$ if $\overline{s_2}(s_1) > q$ and $\overline{s_2}(s_1 + 1) = 0$ if $\overline{s_2}(s_1) \leq q$; ii) $A^*(s_1, s_2) = \overline{A}(s_1, s_2) \ \forall s_2 \geq \overline{s_2}(s_1)$ and $A^*(s_1, s_2) = \overrightarrow{0} \ \forall s_2 < \overline{s_2}(s_1)$.*

When $\omega = \infty$ and $\frac{q}{r}$ is integer, the optimal threshold for the shipment action starts at state $(0, \overline{s_2}(0))$, and the s_2 component of the threshold decreases in a step-wise fashion with a step-size of $\frac{q}{r}$ as s_1 increases. This creates a linearly decreasing stair-shaped threshold; therefore, we refer to these policies as linear-stepwise threshold policies. For instance, Figure 2 illustrates the linear-staircase threshold-type optimal policy for a particular problem instance with $\omega = \infty$. In this figure and those in Section 3.3, W and

13

S indicate the (s_1, s_2) combinations for which the optimal action is to *WAIT* and *SHIP*, respectively. The double-black line represents the optimal staircase threshold policies. In the optimal policy in Figure 2, $\overline{s_2}(0) = 17$ which is sufficient to define the optimal policy under the aforementioned conditions.

The volume of a typical truck is about 100 m^3 (cubic meters). If we use 0.1 m^3 as the size of each unit load, the range of s_i values may be in thousands, which requires manipulation of transition probability matrices for millions of states. The characterization in Theorem 3 makes such real-life size problem instances tractable: Any feasible linear-staircase threshold-type policy can be specified by a threshold s_2 value beyond which the policy requires a shipment action for $s_1 = 0$. Therefore, when $\frac{q}{r}$ is an integer, the uncapacitated model can be solved efficiently by identifying the border of the optimal linear-staircase threshold for $s_1 = 0$, i.e., $\overline{s_2}(0)$ as defined in Theorem 3. However, when $\omega < \infty$ or q is not an integer, the optimal policy may no longer be linear-stepwise; thus, the whole shipment threshold should be specified.

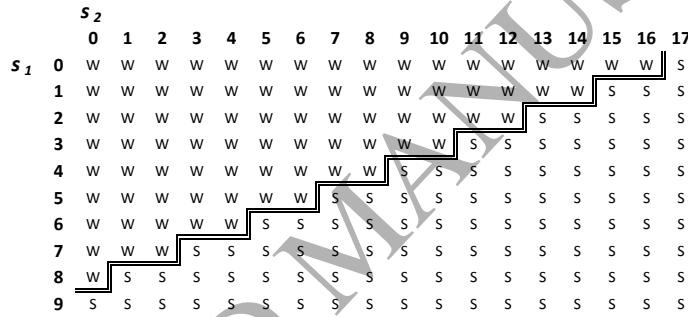| s_2	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
s_1 0	w	w	w	w	w	w	w	w	w	w	w	w	w	w	w	w	w	S
1	w	w	w	w	w	w	w	w	w	w	w	w	w	w	w	S	S	S
2	w	w	w	w	w	w	w	w	w	w	w	w	w	S	S	S	S	S
3	w	w	w	w	w	w	w	w	w	w	w	S	S	S	S	S	S	S
4	w	w	w	w	w	w	w	w	w	S	S	S	S	S	S	S	S	S
5	w	w	w	w	w	w	w	S	S	S	S	S	S	S	S	S	S	S
6	w	w	w	w	w	S	S	S	S	S	S	S	S	S	S	S	S	S
7	w	w	w	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
8	w	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
9	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S

Figure 2: The Optimal Policy for the Problem Instance with $K = 15$, $c_1 = 1$, $c_2 = 0.5$, $\lambda_1 = 1$, $\lambda_2 = 3$, $\alpha = 0.01$, $d_1(1) = d_2(1) = 1$, $\omega = \infty$

The optimal threshold policies exhibit other interesting patterns observed in our numerical experiments. For example, the optimal thresholds are lower for the capacitated version of a problem (i.e., $\omega < \infty$) compared to the uncapacitated version. In addition, for $\omega = \infty$, when one order-size distribution is stochastically larger than another, then $\overline{s_2}(0)$ for the former distribution is greater. Specifically, consider two particular sets of distributions (d_1, d_2) and (d'_1, d'_2), where the first set is stochastically larger than the second one (i.e., $d_1 >_{st} d'_1$ and $d_2 >_{st} d'_2$). Then, $\overline{s_2}(0)$ for (d_1, d_2) is greater than that for (d'_1, d'_2) . For example, let us consider $d_1(2) = d_2(2) = 1$ and $d'_1(1) = d'_2(1) = 1$. The order accumulation is twice as fast in the case of the former distribution. Therefore, under the same threshold policy, a system with distribution d_i incurs higher shipment

14

cost and lower holding cost compared to a system with distribution d_i'. Thus, $\overline{s_2}(0)$ should be greater in the case of the former distribution to have a balance between holding and shipment costs, similar to that established by the optimal policy for the latter distribution. In order to appreciate the customized solution approach proposed in Section 3.4, it is paramount to visualize the aforementioned properties of the optimal policies.

3.3. Hypothetical Examples

Let us now solve a set of hypothetical problems to verify and illustrate the preceding structural properties. In these problem instances, we set $c_1 = 1$, $\lambda_1 = 1$, and $d_1(j) = d_2(j)$ for $j \in \{1, 2\}$. In these experiments, we consider low and high set-up costs (i.e., $K \in \{5, 15\}$); and low, medium, and high vehicle capacity (i.e., $\omega \in \{7, 20, \infty\}$) and arrival rate scenarios (i.e., $\lambda_2 \in \{3, 6, 10\}$). We also consider five holding cost levels (i.e., $c_2 \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$), and three order-size distributions: $(d_i(1), d_i(2)) = (1, 0)$, $(d_i(1), d_i(2)) = (0.7, 0.3)$, and $(d_i(1), d_i(2)) = (0.3, 0.7)$ for $i \in \{1, 2\}$. We refer to these distributions as the No Skewness, Low Skewness, and High Skewness scenarios, respectively, based on their left-skewness levels. According to the usual stochastic order, $(0.3, 0.7) \succ (0.7, 0.3) \succ (1, 0)$. Thus, we have solved 270 hypothetical problem instances.

Figures 2-4 show the optimal policies for the uncapacitated and capacitated versions of three problem instances. In these figures (and in the rest of the paper), the red line shows the vehicle capacity. For the example in Figure 2, $\frac{c_1}{c_2} = 2$. Therefore, the optimal threshold for shipment decisions is linear staircase with $\overline{s_2}(0) = 17$ for the uncapacitated version of the problem. Starting from state $(0, 17)$, the optimal threshold increases by one unit in s_1 for each two-unit decrease in s_2. In this example, the optimal threshold for the capacitated version with $\omega = 20$ is the same as that of the uncapacitated version, because vehicle capacity is large compared to the maximum load to be carried under the optimal policy of the uncapacitated version (i.e., 17 units). We observe that this trend generally holds when $\frac{c_1}{c_2}$ is integer and ω is large enough.

Figure 3 illustrates the optimal policies of another example for which $\frac{c_1}{c_2} = 10$. For the uncapacitated version of the problem with no skewness in the order-size distribution, the optimal threshold policy is linear staircase with $\overline{s_2}(0) = 33$ (see Figure 3a). The optimal policy is still linear staircase for the same problem instance when the order-size distribution becomes highly skewed; however, $\overline{s_2}$ is greater in this case (i.e., the threshold starts at state $(0, 41)$ in Figure 3b).

For the capacitated version of the problem with $\omega = 20$, the optimal policy is not perfectly linear staircase, but still is a staircase policy. A staircase threshold policy has

15

a shipment threshold whose s_2 level decreases in a step-wise manner while s_1 increases, and has a stair-like structure. However, the step-lenghts of a staircase threshold policy may vary in s_1, i.e., $\overline{s_2}(0) - \overline{s_2}(1) = 4$, $\overline{s_2}(1) - \overline{s_2}(2) = 6$, $\overline{s_2}(2) - \overline{s_2}(3) = 10$ in Figure 3c. These differences would be constant in a linear staircase threshold-type policy. The difference between the optimal policies in Figures 3a and 3c is because the maximum load in the uncapacitated version (33 units) is significantly larger than the vehicle capacity (20 units). However, the two optimal thresholds shown in Figures 3a and 3c are identical when $s_2 < 19$. This result implies that for particular cases, the solution of the uncapacitated version may be used to derive a good initial solution to solve the capacitated version with policy-iteration or value-iteration methods. Also note that, when $s_1 = 0$, the optimal policy in Figure 3c suggests to continue consolidation, even though the consolidated load exceeds the vehicle capacity. On the other hand, for $s_1 > 0$, the optimal policy ships the consolidated load before reaching the capacity limit. Actually, in most cases, the optimal threshold policies derived from the MDP model are significantly different than the policy of initiating a shipment whenever the total consolidated load is equal to or larger than the vehicle capacity. Nor do the optimal policies specify a fixed threshold for the total size of the consolidated load to be dispatched.

(a) Uncapacitated Solution

(b) Uncapacitated Solution under High Skewness
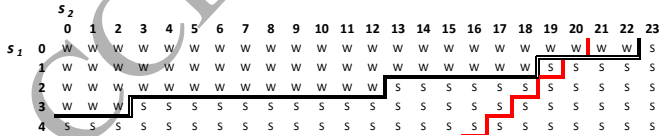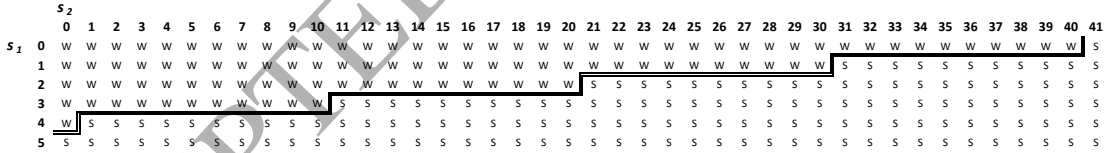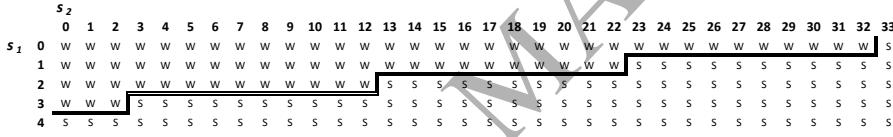

(c) Capacitated Solution

Figure 3: The Optimal Policies for the Uncapacitated and Capacitated Problem Instances with $\omega = 20$, $K = 5$, $c_1 = 1$, $c_2 = 0.1$, $\lambda_1 = 1$, $\lambda_2 = 3$, $\alpha = 0.01$, $d_1(1) = d_2(1) = 1$.
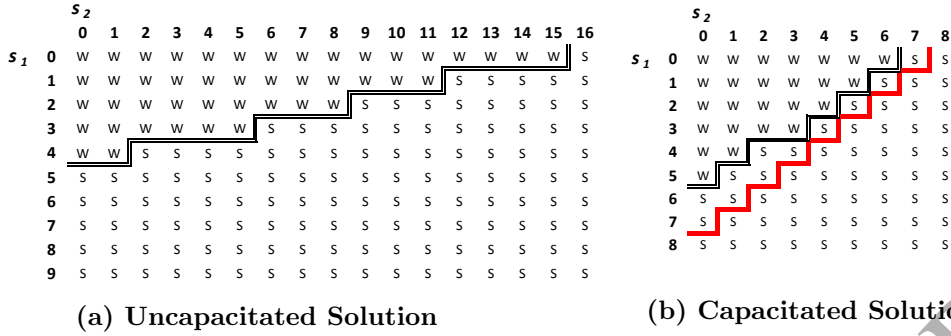
16

Figure 4a: Uncapacitated Solution (rows $s_1 = 0..9$, columns $s_2 = 0..16$)

s_1 \ s_2	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
0	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	S
1	W	W	W	W	W	W	W	W	W	W	W	W	S	S	S	S	S
2	W	W	W	W	W	W	W	W	W	S	S	S	S	S	S	S	S
3	W	W	W	W	W	W	W	S	S	S	S	S	S	S	S	S	S
4	W	W	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
5	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
6	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
7	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
8	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S
9	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S	S

(a) Uncapacitated Solution

Figure 4b: Capacitated Solution (rows $s_1 = 0..8$, columns $s_2 = 0..8$)

s_1 \ s_2	0	1	2	3	4	5	6	7	8
0	W	W	W	W	W	W	W	S	S
1	W	W	W	W	W	W	S	S	S
2	W	W	W	W	W	S	S	S	S
3	W	W	W	W	W	S	S	S	S
4	W	W	S	S	S	S	S	S	S
5	W	S	S	S	S	S	S	S	S
6	S	S	S	S	S	S	S	S	S
7	S	S	S	S	S	S	S	S	S
8	S	S	S	S	S	S	S	S	S

(b) Capacitated Solution

Figure 4: **The Optimal Policies for the Uncapacitated and Capacitated Problem Instances with** $\omega = 7$, $K = 5$, $c_1 = 1$, $c_2 = 0.3$, $\lambda_1 = 1$, $\lambda_2 = 3$, $\alpha = 0.01$, $d_1(1) = d_2(1) = 0.7$, $d_1(2) = d_2(2) = 0.3$

Figure 4 illustrates the optimal policies for a problem instance with a non-integer $\frac{c_1}{c_2}$ ratio (i.e., $r = 3$ and $q = 10$). The optimal threshold for the uncapacitated version of the problem is staircase non-increasing but not linear, e.g., the decrease in s_2 until the next increase in s_1 starts with four units and continues with 3 units. The vehicle capacity in this example ($\omega = 7$) is much lower than the maximum consolidated load to be shipped under the optimal policy of the uncapacitated version (16 units). Thus, the optimal policies in Figures 4a and 4b visibly differ.

Figure 5 shows the average percentage increase (API) in the minimum value functions of the capacitated problem instances (V_C) compared to those of the uncapacitated ones (V_U). We define the percentage increase in each problem instance as $\frac{(V_C - V_U)}{V_U} \times 100\%$. Figure 5 shows that API is very low for the instances with $\omega = 20$ because the vehicle capacity is large compared to the total number of shipment orders on the optimal load-dispatching threshold in most of the uncapacitated problem instances. This result supports our earlier observation: When the capacity of the vehicle is large enough, the optimal policies for the uncapacitated problems generally provide good approximations to those of the capacitated problems. As expected, API increases as the holding cost of regular orders (c_2) decreases, and as the interarrival rate of regular orders (λ_2) and shipment cost (K) increase. This is reasonable because i) when c_2 decreases and K increases, holding inventory becomes cheaper compared to shipping; thus, the decision maker would keep more orders in inventory between shipments. As a result, the likelihood of exceeding the vehicle capacity increases. Increasing λ_2 has a similar effect as it leads to an increase in shipment order accumulation. In addition, stochastically larger order-size distribution (i.e., greater skewness in our examples) is also associated with

17

larger API.

3.4. Customized Solution Approach

Deriving the optimal solution by solving Equations (5) or (6) can be challenging for real-life problems which may have thousands or millions of states. The conventional policy iteration and value iteration algorithms may not work efficiently on such large problem instances, due to requiring computationally expensive matrix inversion operations or having slow convergence. Therefore, we develop a customized solution approach to derive the optimal policies for the uncapacitated problem instances when $\frac{c_1}{c_2} = \frac{q}{r}$ is integer. The proposed algorithm efficiently evaluates possible linear-staircase thresholds and determines the optimal one when $\omega = \infty$, based on two findings.

18
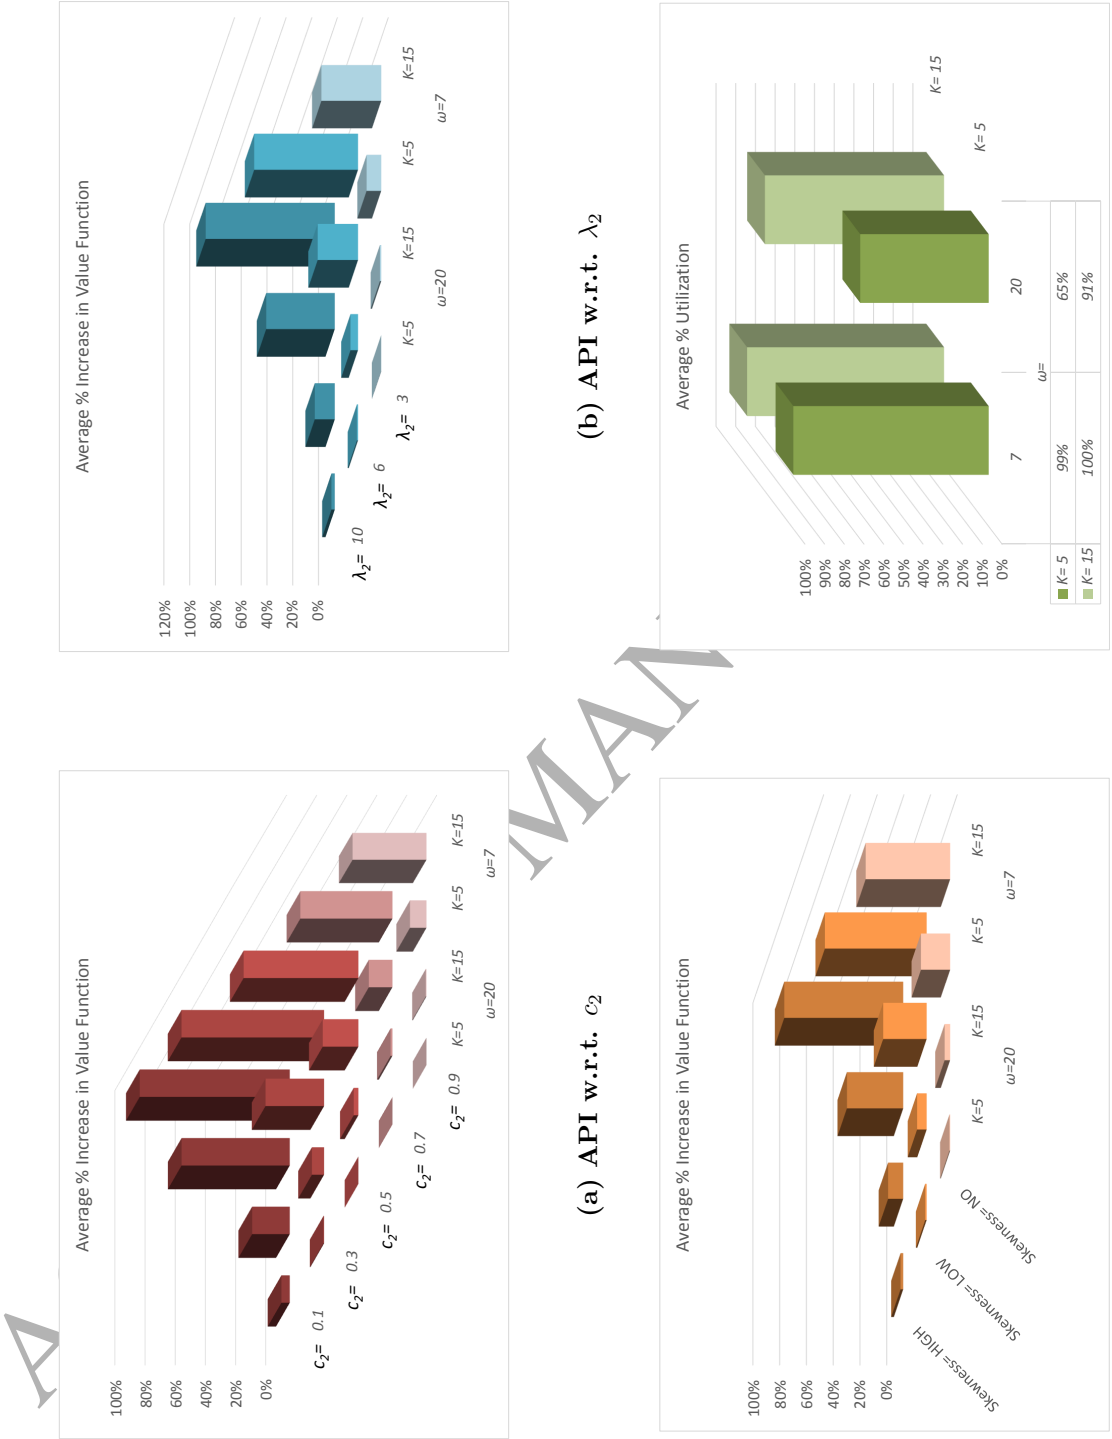
(a) **API w.r.t.** c_2

(b) **API w.r.t.** λ_2

(c) **API w.r.t Order-Size Distributions**

(d) **Average Percentage Capacity Utilization**

Figure 5: Results of Hypothetical Examples

First, since each linear-staircase threshold can be specified by its border at $s_1 = 0$, the optimal threshold can be found by evaluating a finite number of such linear-staircase thresholds. We prove this by deriving a lower (s_2^{LB}) and upper bound (s_2^{UB}) for $\overline{s_2}(0)$ of the optimal policy as described in Lemmas 3 and 4.

Lemma 3. *When* $\omega = \infty$, $\overline{s_2}(0) \leq s_2^{UB} = \lceil \frac{K(\alpha+\lambda)}{c_2} \rceil$.

Lemma 4. *When* $\omega = \infty$, $d_1(1) = d_2(1) = 1$, *and* $A^*(e_1) = A^*(e_2) = \overrightarrow{0}$, *then* $\overline{s_2}(0) \geq$ $s_2^{LB} = \frac{1}{c_2}\left(\alpha K + \frac{c_1\lambda_1}{\alpha+\lambda} + \frac{c_2\lambda_2}{\alpha+\lambda}\right)$.

Note that the condition $A^*(e_1) = A^*(e_2) = \overrightarrow{0}$ in Lemma 4 generally holds unless K is very small compared to either c_1 or c_2. This condition can be easily verified by checking whether the linear-staircase threshold policies with a threshold border at state $(0, q + 1)$ outperform those with lower thresholds. In addition, our numerical experiments illustrated that this lower bound applies to the problems with batch arrivals (i.e., $d_1(1) < 1$ and $d_2(1) < 1$).

Actually, the lower bound condition in Lemma 4 can be extended to the batch-arrival case by generalizing the assumptions accordingly. Let the expected size of a Type i shipment order be D_i. In case of batch arrivals and $\omega = \infty$, the lower bound s_2^{LB} is equal to $\frac{1}{c_2}\left(\alpha K + \frac{c_1\lambda_1}{\alpha+\lambda}D_1 + \frac{c_2\lambda_2}{\alpha+\lambda}D_2\right)$ if $A^*(Ne_1) = A^*(Me_2) = \overrightarrow{0}$. Although it is a stronger condition than $A^*(e_1) = A^*(e_2) = \overrightarrow{0}$, $A^*(Ne_1) = A^*(Me_2) = \overrightarrow{0}$ may be a reasonable assumption when N and M are small enough or/and K is large enough. This assumption can be verified in a manner similar to that of the preceding paragraph, i.e., by checking whether the linear-staircase threshold policy with textcolorreda threshold border at state $(0, \max(M + 1, qN + 1))$ outperforms those with lower thresholds. The proof of the generalized lower bound is given in Appendix C.

Secondly, the performance of each threshold policy can be efficiently evaluated using a different approach. Conventionally, each linear-staircase threshold policy can be represented as a Markov chain. However, evaluating those Markov chains may again require inverting a large matrix. As an alternative, we analytically express $V^{\pi(s_2)}(0,0)$, the expected total discounted cost of a particular linear-staircase threshold policy with a threshold border at state $(0, s_2)$ for any integer $s_2 \in [s_2^{LB}, s_2^{UB}]$ given the initial state $(0,0)$, as a recursive equation with a single unknown variable. Then, the optimal policy can be identified by deriving $\overline{s_2}(0) = argmin_{s_2^{LB} \leq s_2 \leq s_2^{UB}}\{V^{\pi(s_2)}(0,0)\}$. Details of the aforementioned recursive cost function are available in Appendix D. Section 4.2 explains how this customized solution approach is utilized within the overall solution methodology.

20

4. Computational Analysis

In order to assess its potential to improve industrial systems, we applied the proposed model to the cases of two logistics firms operating in Turkey, using real data. EKOL Logistics is a 3PL serving manufacturers on defined milk-run routes. We focus on EKOL's services for automotive manufacturers in the central and western Anatolia region, and determine when a truck should be dispatched on a route to collect the goods ordered until then, and carry them to consignees. UPS Turkey is a courier that delivers parcels using their own trucks, as well as hired vehicles. Parcels dropped at various UPS stores or collected from customers are transferred between hubs and then distributed to different end points. For UPS Turkey, we focus on the parcel traffic between the two main hubs, İstanbul and Ankara, and determine when the parcels accumulated in İstanbul should be sent to Ankara.

4.1. Data and Parameter Estimation

The data from EKOL include detailed records of automotive part shipments on milk-run routes during the month of April 2015. The data consist of 2,490 shipment records excluding those for the returned empty containers. These records specify information such as shipment order size, truck details, time of the shipment, and origin/destination. Because 60% of these records concern shipments from a set of suppliers to two large manufacturers, we concentrate on planning the routes involving those manufacturers. Note that the remaining 40% of the records do not share any common shipment with those for the two manufacturers.

For each manufacturer, we specify the suppliers on the same shipment route using a community detection algorithm. We calculate the times that items from supplier i are shipped in the same vehicle with those of supplier j (l_{ij}). Considering these suppliers as vertices and l_{ij} as the edge weight, we establish a weighted graph structure for each manufacturer as shown in Figures 6a and 6c. These graphs represent suppliers as nodes in different colors, and connect two nodes if they are ever shipped together. Nodes that are well-connected with each other are assigned the same color. The matrices next to each graph indicate how strongly each node is connected with the others, i.e., the size of the sign in each matrix cell is proportional to the number of times that items of supplier i and j are shipped together. Figures 6a and 6c illustrate that for each manufacturer, there are at least three supplier groups whose items are regularly shipped together. However, some supplier groups are well-connected with other groups. Therefore, we employ a community detection algorithm, i.e., the Fast Greedy approach of Clauset et al. (2004) using R version 3.2.4, to divide suppliers of each manufacturer into groups/communities

21

that maximize the modularity of the weighted graph structure. This way, we obtain supplier groups that are well-connected within themselves, and sparsely connected with the other supplier groups.

Figures 6b and 6d show the respective supplier groups for Manufacturer 1 and 2, as derived by the algorithm. For Manufacturer 1, the algorithm finds two sparsely connected supplier groups. Therefore, we assign two milk-run routes for Manufacturer 1, namely 1A and 1B. For Manufacturer 2, the algorithm identifies three major groups: Group 3 (green nodes in Figure 6c) has shipments in common with the other two, while items from Groups 1 and 2 are not shipped together. We duplicate the suppliers in Group 3 and distribute the shipment records between them, depending on the group with which they are shipped (1 or 2). This way, we identify two distinct milk-run routes for Manufacturer 2 as well, namely 2A and 2B. The reordered matrices in Figures 6b and 6d illustrate that the suppliers in each new supplier group (1A, 1B, 2A, and 2B) rarely have items that are shipped together with those of other supplier groups.

(a) Manufacturer 1-Initial

(b) Manufacturer 1-Final

(c) Manufacturer 2-Initial

(d) Manufacturer 2-Final
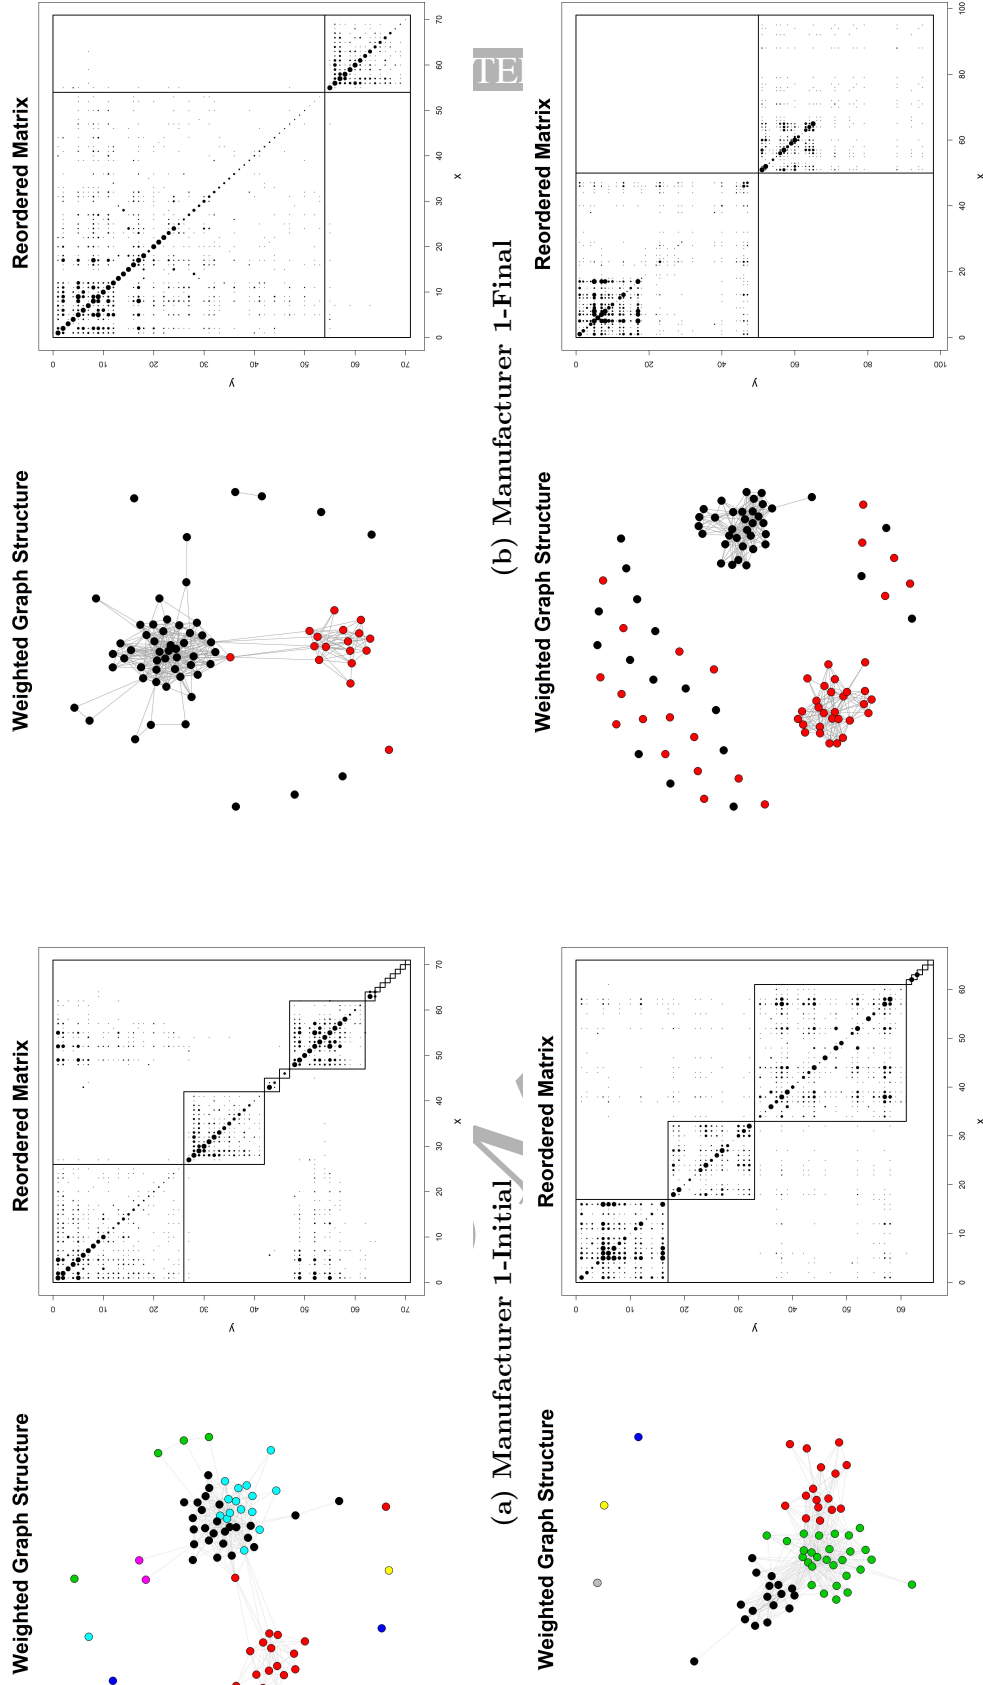
Figure 6: Supplier Groups Identified by the Community Detection Algorithm

Our data on when the parts from the suppliers were ready for pickup are not sufficiently detailed for statistical testing to derive the interarrival times. However, the superposition of numerous independent arrivals from many independent suppliers can be approximated as a Poisson process by the central limit theorem. We thus assume exponential interarrival times. For the size distribution of shipment orders, we use a discretization approach similar to that of Higginson and Bookbinder (1995). That is, we derive empirical distributions for $d_1(n)$ and $d_2(m)$ by categorizing the data on shipment order sizes into discrete groups using a unit load size of 4.5 m^3, where a great proportion of the actual loads are larger than this value.

The majority of shipments are carried out in trucks of volume 90 m^3; therefore, we take ω=20 ($90m^3/4.5m^3$) for the EKOL case. We are aware that this approach neglects other dimensions of shipment orders; however, applying the truck capacity constraint based on a total load volume is reasonable given that EKOL's average truck utilization is less than 65%. Then, depending on when manufacturers need them, orders are defined as EARLY, OK, or LATE. We treat LATE as *expedited* and the other two as *regular* items. We estimate the holding cost pairs (c_1, c_2) based on average truck capacity utilization (ATCU). For this purpose, we first define $\rho = \frac{c_1}{c_2}$; and set $\rho = 3$ and $\rho = 6$ as high and low holding cost ratios, respectively. Then, we performed a search over (c_1, c_2) on the capacitated model, and found the values for which the expected vehicle utilization of the optimal policy derived from the capacitated model (ω=20) is equal to the ATCU calculated from EKOL's data. We call the resulting (c_1, c_2) values as the base case holding costs. We derive the K values based on truck rental costs used in actual practice. Normalized on K, the average c_2 values for all routes are $0.0675K$ and $0.0558K$ for $\rho = 3$ and $\rho = 6$, respectively. We also consider 150%, 75%, 50% and 25% of the base case (c_1, c_2) values to measure sensitivity of the results to the holding costs.

The data from UPS Turkey include the records of 28,577 parcels carried between the main hubs of Istanbul and Ankara in December of 2015. The records belong to 128 UPS stores and comprise detailed parcel information including each parcel type, weight, dimensions, price, receipt/arrival and delivery times, delivery location, and delivery status (late, on-time) as well as truck information (volume capacity, cost of dispatch, vehicle license number). The data were cleansed by excluding the repeated records and records on returns. In addition, the records from the same customer at a specific store within a small time interval (less than 5 minutes) were combined into a batch order arrival. The remaining data consist of 14,284 parcel records. The interarrival times are calculated and fit to an exponential distribution. The goodness of fit tests for exponential interarrival times are conducted using Minitab V.15.1 on the data from seven stores that

receive more than 50% of the parcels. The associated *p-values* range between 9.4% and 98.1%, implying that it is reasonable to assume exponential interarrival times. The proportion of expedited and regular shipment orders are calculated from parcel-type information. We included the parcels with tight delivery time promise (e.g. under the risk of being late) in the expedited order category.

UPS generally uses two types of trucks whose volumes are 48 and 100 m^3. All shipments performed by larger trucks had a total load volume of less than 48 m^3. In fact, the ATCU is calculated as 45.21%, even assuming that all shipments are done with the smaller truck. Assuming a unit-load size of $0.1m^3$, we applied the proposed model to the case of UPS Turkey with a truck capacity $\omega=480$ ($48m^3/0.1m^3$). This unit-load size limits the possible shipment-order sizes to three for each shipment type, and enables us to solve this problem. We set $\frac{c_1}{c_2} = 5$ because UPS Turkey charges five times as much for carrying expedited orders compared to regular orders. Then, we derive the (c_1, c_2) pairs for which the ATCU calculated from the data of UPS Turkey approximately matches the expected vehicle utilization achieved when the optimal quantity-based policy for the uncapacitated case is applied.

4.2. Solution Methodology

We solve the shipment consolidation and capacity rationing problems for the cases of EKOL Logistics and UPS Turkey employing the solution procedure depicted in Figure 7. This procedure first uses the customized solution approach described in Section 3.4 to derive good initial solutions, and then apply a value iteration algorithm to derive the optimal or good solutions for the capacitated problem instances. The value iteration algorithm is described in the following pseudo-code. We denote the value function derived at the k^{th} iteration of the algorithm based on Equation (7) as $V_k(S)$, where V_k is the vector form of the value function.

Algorithm 1. *Pseudo-code for Value-Iteration Algorithm*

1: **Set** V_0, ϵ, β
2: **while** $\|V_k - V_{k-1}\| > \epsilon\frac{(1-\beta)}{(2\beta)}$ **do**
3:    **function** VALUEITERATION(V_k)
4:      **Solve** V_k ▷ *using Equation (7)*
5:    **end function**
6: **end while**
7: **return** $\{V, \pi^*\}$

25

$$V_k(S) = min\Big\{\frac{1}{\alpha+\lambda}CS^{tr} + \sum_{n=1}^{N}\frac{d_1(n)\lambda_1}{\alpha+\lambda}V_{k-1}(S+ne_1) + \sum_{m=1}^{M}\frac{d_2(m)\lambda_2}{\alpha+\lambda}V_{k-1}(S+me_2),$$

$$K + \sum_{n=1}^{N}\frac{d_1(n)\lambda_1}{\alpha+\lambda}V_{k-1}(ne_1) + \sum_{m=1}^{M}\frac{d_2(m)\lambda_2}{\alpha+\lambda}V_{k-1}(me_2)\Big\} \tag{7}$$

Phase 1 of the solution procedure in Figure 7 employs the customized solution approach described in Section 3.4 to find the optimal solution for the uncapacitated version of the problems. That is, this procedure evaluates all linear-staircase threshold policies specified by possible $\overline{s_2}(0)$ values in the range of $[s_2^{LB}, s_2^{UB}]$ and selects the best one. The customized solution approach is developed for integer $\frac{c_1}{c_2}$ ratio. Therefore, if $\frac{c_1}{c_2}$ is not integer, it is rounded and the customized approach is applied. Then starting from the solution of the uncapacitated problem with rounded $\frac{c_1}{c_2}$, a value iteration algorithm is run to find the actual optimal solution of the uncapacitated problem.

Phase 2 of the solution procedure feeds the policy derived in Phase 1 to a value iteration algorithm that considers the vehicle capacity. It is possible that the optimal policy will no longer be linear-staircase, but a staircase one after applying the capacity restriction, as explained in Figure 3. Having a good initial solution, the value iteration algorithm converges to an optimal solution much faster. This approach is feasible up to a certain problem size. If the problem size is very large (e.g., the case of UPS Turkey), we apply the value iteration algorithm by enforcing linear-staircase threshold policies.

Because linear-staircase policies may not be optimal when $\omega < \infty$ (see Figure 3c for an example), we check whether solution of the value iteration algorithm can be improved via a greedy neighbourhood search. Starting from the policy found by the value iteration step and $i = 0$, the neighborhood search myopically checks whether reducing/increasing $\overline{s_2}(i)$ improves the performance of the current policy or not. When the improvement stops for i, the algorithm continues by increasing i by one unit. The algorithm terminates when $\overline{s_2}(i) = 0$ and increasing $\overline{s_2}(i)$ by one unit does not improve the system performance. The optimality of the final solution is verified via a policy-improvement step (as done in a conventional policy iteration algorithm). Finally, the quantity-based (optimal) policies derived by this procedure are compared via simulation with the time policies practiced by EKOL Logistics and UPS Turkey. These time policies are mainly periodic policies, whose schedules are derived using the shipment records in the data from EKOL Logistics and UPS Turkey.

The aforementioned solution approach is coded in MATLAB and run using a PC with Intel Pentium 4 Processor and 4GB RAM. The time policies practiced by EKOL
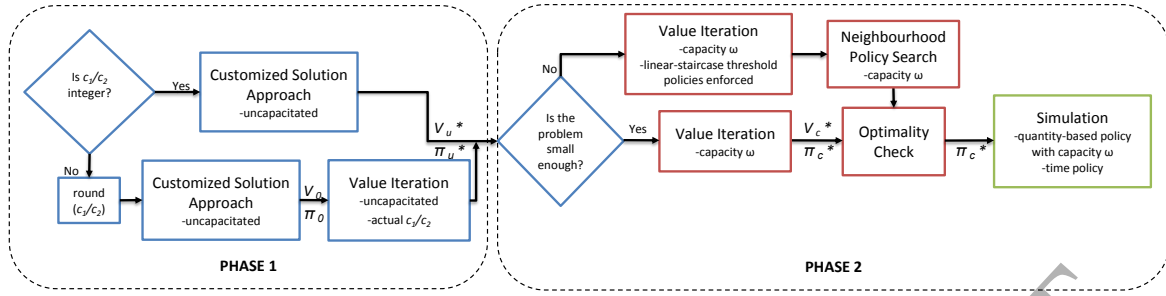
26

Figure 7: Solution Methodology

Logistics and UPS Turkey are tested in ARENA V.13.5 with 50 replications. Each replication is run for one year, with a warm-up period of 2 months.

4.3. Numerical Results

In EKOL's case, we solved 40 problem instances in total (4 routes x 2 ρ levels x 5 holding-cost scales). Figure 8 illustrates the optimal shipment policies for one of the EKOL problem instances (milk-run route 2B with $\alpha = 0.01$, c_2/K=0.0167, and base-case holding cost) under both the vehicle capacity scenarios. Note that, while the total number of shipment orders in the optimal load-dispatching threshold is less than or equal to the vehicle capacity (20 units), the optimal policy for the capacitated case has a lower threshold, i.e., $\overline{s_2}(0) = 13$ in Figure 8b. This is because the sizes of shipment orders vary significantly in EKOL's case. If the optimal policy in Figure 8a were employed when $\omega = 20$, the consolidated load may exceed the vehicle capacity at the next order arrival, and incur additional holding cost until the following shipment. The lower optimal threshold in Figure 8b aims to eliminate such possibilities to an extent. For instance, for the un-capacitated problem, the optimal decision is to wait at state (1,15) as indicated in Figure 8a. If the next order is type 2 with a size of 7, the system state moves to (1,22), whose consolidated load exceeds the capacity of a truck (20 units). However, the lower optimal shipment threshold for the capacitated problem shown in Figure 8b dispatches the consolidated load when the system state reaches state (1,15), and prevents such a possibility.

Figure 9 shows the mean of the expected shipment, holding and total costs associated with the optimal quantity-based policies and the time policies practiced by EKOL Logistics for all routes and ρ values. The mean is taken over the holding-cost scenarios. On average, the optimal policies reduce the expected total cost by 48%. The percentage reduction varies between 32% and 62%. Although the optimal policies reduce both shipment and holding costs in most cases, the reduction in holding cost is more visible

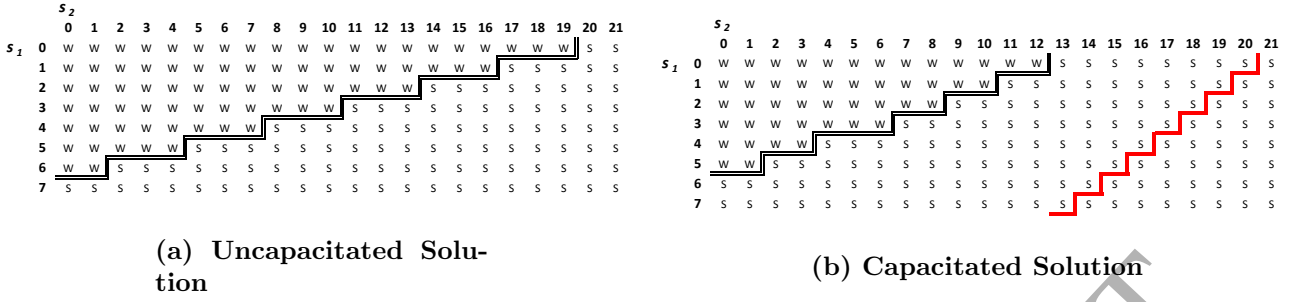27

(a) Uncapacitated Solution

(b) Capacitated Solution

Figure 8: Optimal Solutions for EKOL's Problem

in Figure 9. We also observe that the expected times between consecutive dispatches for the time policies and for the optimal quantity-based policies are similar. However, the variance of the time between consecutive dispatches is greater in the latter group, e.g., the average coefficient of variation (CV) for time between consecutive dispatches is 0.74 for the optimal quantity-based policies.
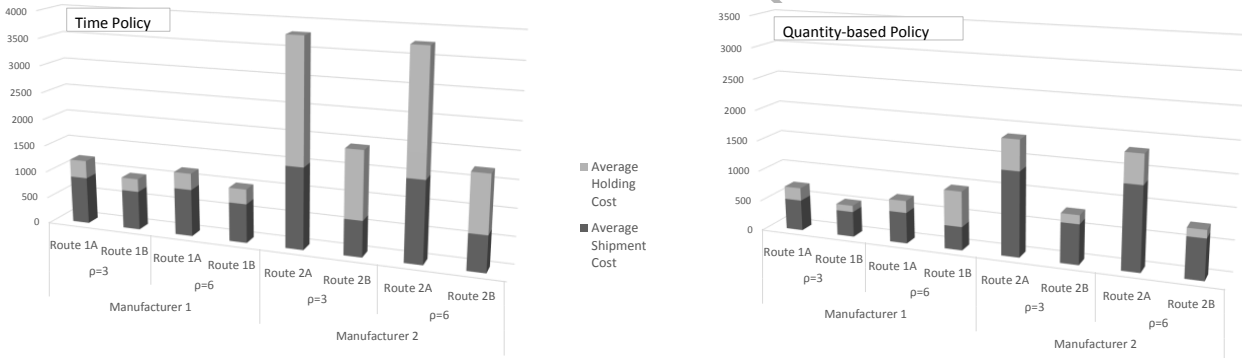

Figure 9: Performances of the Optimal Quantity-Based Policies and Time Policies for EKOL's Problem Instances

For the case of UPS Turkey, we set the unit load as 0.1 m^3; therefore, we have the truck size $\omega = 480$ units. After deriving the optimal policies for this unit-load size, we evaluated their performance using the ARENA simulation model for a much smaller unit load (0.005 m^3). The smaller unit-load size allows analysis of the optimal policy in a more granular and realistic setting with 100 possible order sizes and a maximum volume of 9,600 units for the load dispatched. Naturally, the optimal policy for a unit load of 0.1 m^3 can be applied to a more granular setting in the simulation, *starting from different shipment threshold border (i.e., $(0, \overline{s}_2(0))$.* Therefore, we perform a neighbourhood

search via simulation to derive the best $\overline{s_2}(0)$ option for the more granular setting. We compare this implementation of the optimal quantity-based policies with two time-based policies: a time policy and a modified time policy. For the time policy, the interarrival time between two consecutive dispatches is equal to the mean time between dispatches reported in the UPS data (i.e., 4.38 hours). For the modified time policy, that interval is equal to the mean time between consecutive dispatches under the optimal quantity-based policy (3.966 hours).

Table 1 provides the comparison of these three policies. The quantity-based policy outperforms the two time-based policies by reducing the total cost by up to 3.2%. In this case, the quantity-based policy indicates slightly more frequent shipments than the time policy. However, the increase in transportation cost is compensated by a reduction in holding cost. In addition, the optimal quantity policies perform slightly better than the time policies in terms of proportion of timely-shipped orders and average lateness among the late orders. This implies that the proposed optimal quantity policies may reduce costs without additional violation of promised due-dates compared to the current practice. These results provide two important insights. First, compared to EKOL's case, the improvements achieved by the quantity-based policy from the proposed model is limited in the case of UPS Turkey. This is mainly because the shipment order sizes for UPS Turkey are very small compared to the vehicle capacity; therefore, the variation of the time until the next shipment under the quantity-based policy is limited compared to that of EKOL. Note that CV of the time between consecutive dispatches for UPS Turkey is 0.108, much smaller than EKOL's value of 0.74. Second, the performance of the modified time policy is better than the time policy. This implies that a significant portion of the benefits achieved by employing the optimal quantity-based policies may also be obtained by adjusting the shipment frequency of the practiced time policies in cases similar to those of UPS Turkey.

Holding cost is an indirect cost representing the negative effects of lateness in transporting shipment orders, such as increased material handling/storage costs and reduced customer satisfaction. Therefore, it may be desirable to derive alternative quantity policies reducing both holding and shipment costs in cases like that of UPS-Turkey. The proposed MDP model can be used to identify such policies by systematically adjusting the holding costs to find an alternative policy with a less aggressive shipment schedule. A search protocol for this purpose is presented in Appendix E. Following this search protocol, we found an alternative quantity policy with a $\overline{s_2}(0)$ value that is 8.6% greater than that of the original optimal solution. Compared to the time policy, this alternative quantity policy is associated with i) a similar rate of timely-shipped orders, and ii) 2.75%

29

Table 1: Cost Reductions in Lateness Measures and Costs Achieved by the Quantity-Based Policies Compared to the Time and Modified-Time Policies for UPS Turkey.

Quantity-Based Policy over:	Proportion of Timely-Shipped Orders	Average Lateness of Late Orders	Shipment Cost	Holding Cost	Total Cost
Time Policy					
Expedited Orders	2.2%	-0.7%	6.26%	-10.95%	-3.20%
Regular Orders	0.9%	-6.4%			
Modified-Time Policy					
Expedited Orders	0.1%	0.8%	-3.82%	-1.22%	-2.53%
Regular Orders	-0.3%	35.3%			

Improvements are indicated with positive values in Column 1 and with negative values in the other columns. The proportion of timely shipped orders is above 95% for both order types under the optimal quantity policies.

and 2.18% less total and shipment costs, respectively.

5. Summary and Conclusions

The goal of shipment consolidation is to attain economies of scale, spreading the fixed transportation cost over a greater number of orders. However, the total number of orders in the consolidated load may not be as great as hoped. That is either because of the limited vehicle capacity, or the degrading of customer service by the possibly-long lead time before delivery for those customers whose orders were first to arrive.

In this paper, we have dealt with the latter difficulty by prioritizing the orders. Our continuous-time MDP model considers two classes of orders, the first of which receives greater consideration than the second in making up a load for dispatch. We thus "ration" the capacity of the transportation vehicle, in allocating the volume of the truck between the expedited (Type 1) and regular (Type 2) orders to minimize the expected total discounted cost incurred over an infinite horizon. The cost structure includes the fixed cost per vehicle dispatched, plus the holding costs, $c_1 > c_2$, whose values reflect the priority of the Type 1 orders. To enable efficient solution of this problem, we characterize the optimal policies, which are of control-limit type with linear-staircase thresholds for the case of an uncapacitated vehicle, i.e., $\omega = \infty$ (see Sec. 3.2, Theorem 3). These thresholds refer to how many unit-size orders of each type to consolidate before dispatching a load. Note that this policy is thus considerably more complicated than deciding to ship when the consolidated load reaches a particular size.

We propose a solution approach (Figure 7) relying on the following two points: 1) the expected discounted total cost of a given linear-staircase threshold policy can be calculated efficiently when $\frac{c_1}{c_2}$ is integer, 2) there are a finite number of such thresholds to search among to find the optimal one. The proposed solution method is employed on the industrial applications for EKOL Logistics and UPS Turkey (Sec. 4.3), which motivated this research.

Our numerical experiments illustrate that the shipment consolidation policies derived for both EKOL and UPS Turkey reduce the total cost compared to the time policies (shipments at fixed intervals) currently used by those firms. The results are more striking for EKOL because their median order size is a significant fraction of the vehicle capacity. In such cases, dynamic control of the shipment decisions becomes very critical in preventing unnecessary transportation and holding costs. For UPS Turkey, there was still an improvement, although smaller than that of EKOL's. This was because the order sizes and their variance are much smaller for UPS Turkey. We also observe that UPS Turkey can improve its cost performance by increasing its shipment frequency to match that of the policy we derived, illustrating another benefit of our modeling approach.

We remark that, although the optimal linear-staircase thresholds can be derived for the case of $\omega = \infty$, they are useful for the capacitated vehicle case as well. The former solution can be used to derive a good initial solution to the value function, when solving the capacitated version via the methods of policy iteration or value iteration. Moreover, there are situations where the optimal thresholds are the same because vehicle capacity is large compared to the total number of shipment orders on the optimal load-dispatching threshold of the uncapacitated version (see Figure 2).

Although we focus on the shipment consolidation and vehicle capacity rationing problem in this paper, the proposed model is applicable to other settings. For instance, our model can be a piece of a larger logistics network model. The shipment loads dispatched to manufacturers by EKOL may not involve end-items, but rather components or sub-assemblies. The supply chain nodes that are further downstream may thus benefit from the shipment consolidation decisions that are derived by our CTMDP model.

In addition, our model is applicable to other batch service problems which involve a set-up and orders (jobs) with different urgency levels. Such problems may exist in medical testing laboratories. For example, pathologists are known to batch-process pathology specimens (e.g., biopsied lesions to be examined for malignancy), and the samples from different patients may have different priorities (Volmar et al. 2013, Hartman 2015). Our model can be employed or extended to determine the optimal time to batch-process such medical tests.

31

Admittedly, our analysis considers only two classes of orders. This is reasonable in the setting of Turkey and other cases where the relevant distances are short. We leave to future research the case of three or more order types. Furthermore, we employ just a single size of truck in this paper; this was justified for our analysis since the load factor (vehicle utilization) was low in the data from EKOL and UPS Turkey. In addition, companies such as UPS are known to use a single type of truck. However, the possibility of multiple types is worthy of additional study. A wider set of actions may then be possible, e.g. *WAIT*; *SHIP* in Type 1 vehicle; *SHIP* in Type 2 vehicle. The potential dispatch quantities could then be more varied.

Moreover, we make particular simplifications in our model: 1) our state space only keeps track of the accumulated expedited and regular orders in unit-order sizes; 2) we use the holding costs to penalize lateness without hard due-date constraints. These simplifications, commonly utilized in the existing literature, are necessary to keep the system state compact for solving large problem instances and benefit from the structural results of existing studies. It is reasonable to have these simplifications in our model, as they do not prevent an accurate representation of our problems dynamics. The former simplification may allow orders from the same customer to be divided and dispatched in separate shipments. Nevertheless, L-Ps will avoid such order-splits when applying the optimal policies, and may incur additional costs. However, our preliminary numerical analysis with the simulation model described in Section 4 shows that the effect of this simplification is limited. For instance, the optimal policy for EKOLs case of Manufacturer-1 Route-A with $\rho = 3$ and the base case costs is associated with a 44% total cost reduction compared to current practice when order-split is allowed. By avoiding the order-split during the implementation of the optimal policy, the L-P sacrifices only 4% of this cost reduction.

Similarly, the latter simplification ignores the promised shipment due dates; however, the holding costs may effectively limit the number of late deliveries, e.g., more than 95% of orders are delivered on-time under the optimal policy in the case of UPS Turkey. We have also measured the combined effect of relaxing both the former and latter simplifications on the performance of the optimal quantity policies in a large numerical experiment (See Appendix F). We have observed that optimal quantity policies derived by the MDP model can improve shipment consolidation practices in terms of cost for various problem settings, even if the assumptions of not explicitly considering due dates and of allowing splitable shipment orders are relaxed simultaneously. However, occasionally, the optimal quantity policies derived by the MDP model may not achieve the desired proportion of timely-shipped orders. In those instances, an alternative quantity policy with a better

32

proportion of timely-shipped orders can be derived by systematically increasing holding costs in the MDP model, and forcing the proposed approach to find a more preferable policy (by sacrificing a portion of the achievable total-cost reduction). Alternatively, these simplifications can be avoided by assuming periodic shipment dispatches whose compositions are determined based on 0-1 Knapsack problems given the exact sizes of accumulated orders. Our additional numerical experiments illustrated that such an approach may work well for particular applications when vehicle-size-to-order-size ratio is small, e.g., the case of UPS Turkey. However, in several other cases like EKOL's, the optimal quantity-based policies reduce the total cost significantly compared to such alternative periodic policies.

Acknowledgements

NOTATION

λ_1, λ_2: Poisson arrival rate for expedited (Type 1) and regular (Type 2) orders.

$\lambda = \lambda_1 + \lambda_2$: Cumulative arrival rate.

$d_1(n)(d_2(m))$: Probability that the size of an arriving expedited (regular) order is n (m) units, where $n \in \{1, 2, ..., N\}$ ($m \in \{1, 2, ..., M\}$).

D_i: Expected size of a Type i shipment order.

K: Fixed cost to dispatch a vehicle.

ω: Vehicle capacity.

c_i: Holding cost, per unit order per unit time, for Type i order.

$C = (c_1, c_2)$: Vector of holding costs.

p: Decision epoch specified by the number of shipment order arrivals so far.

T_p: The random variable representing the time of p^{th} order arrival.

ϵ: A negligibly small value. It refers to allowable optimality gap in Algorithm 1.

α: Continuous discount rate. $\beta = \frac{\lambda}{\alpha+\lambda}$: Discount rate per decision epoch.

$S = (s_1, s_2)$: System state, where s_i denotes the number of Type i orders in the system.

S: Set of states.

33

S_t, A_t: System state and action at time t.

$A(S) = \{a_1, a_2\}$: Action taken in state S, i.e., shipping a_i units of Type i orders.

$\mathbf{A}(S)$: Set of feasible actions in state S.

$I_{[A]}$: Indicator function which is equal to 1 if $a_1 + a_2 > 0$, 0 otherwise.

$\vec{0} = (0,0)$: Vector of zeros. $e_1 = (1,0)$, $e_2 = (0,1)$: Elementary unit vectors.

$p(S'|S, A)$: Transition probability of moving from state $S = (s_1, s_2)$ to state $S' = (s_1', s_2')$ under action $A = (a_1, a_2)$ (just after an order arrival).

$\overline{A}(S) = (a_1, a_2) = (min\{s_1, \omega\}, min\{\omega - a_1, s_2\})$: The action that prioritizes the expedited orders in utilizing vehicle capacity.

$A^*(S)$: Optimal action taken in state S.

$C(S, A)$: One-epoch cost function (immediate cost between two consecutive order arrivals), when state is $S = (s_1, s_2)$ and action is $A = (a_1, a_2)$.

π: A stationary policy. π^*: The optimal stationary policy.

$V(S, A)$: Expected total discounted cost, given that action A is chosen when the initial state is S.

$V(S)$: Minimum total expected discounted cost, given that initial state is S.

V_C (V_U): Optimal value function of the capacitated (uncapacitated) problem.

$\overline{s_2}(0)$: s_2 component of the optimal shipment threshold when $s_1 = 0$.

s_2^{LB} (s_2^{UB}): A lower (upper) bound for $\overline{s_2}(0)$.

$\rho = \frac{c_1}{c_2}$: Cost ratio.

$V_k(S)$: Value function derived in the k^{th} iteration of the value-iteration algorithm.

V_k: Matrix of all value functions derived in the k^{th} iteration of the value-iteration algorithm.

References

References

Alagöz, O., Maillart, L. M., Schaefer, A. J., Roberts, M. S., 2004. The optimal timing of living-donor liver transplantation. Management Science 50 (10), 1420–1430.

Alagöz, O., Maillart, L. M., Schaefer, A. J., Roberts, M. S., 2007. Determining the acceptance of cadaveric livers using an implicit model of the waiting list. Operations Research 55 (1), 24–36.

Bertsekas, D. P., 2001. Dynamic programming and optimal control, vol 2. Belmont, MA, Athena Scientific.

Bookbinder, J. H., Cai, Q., He, Q.-M., 2011. Shipment consolidation by private carrier: The discrete time and discrete quantity case. Stochastic Models 27 (4), 664–686.

Bookbinder, J. H., Higginson, J. K., 2002. Probabilistic modeling of freight consolidation by private carriage. Transportation Research Part E 38 (5), 305–318.

Çetinkaya, S., Bookbinder, J. H., 2003. Stochastic models for the dispatch of consolidated shipments. Transportation Research Part B 37 (8), 747–768.

Çetinkaya, S., Mutlu, F., Wei, B., 2014. On the service performance of alternative shipment consolidation policies. Operations Research Letters 42 (1), 41–47.

Clauset, A., Newman, M. E. J., Moore, C., 2004. Finding community structure in very large networks. Physical Review E 70 (6), 066111.

DallOrto, L. C., Crainic, T. G., Leal, J. E., Powell, W. B., 2006. The single-node dynamic service scheduling and dispatching problem. European Journal of Operational Research 170 (1), 1–23.

de Véricourt, F., Karaesmen, F., Dallery, Y., 2000. Dynamic scheduling in a make-to-stock system: A partial characterization of optimal policies. Operations Research 48 (5), 811–819.

Fadıloğlu, M. M., Bulut, Ö., 2010. A dynamic rationing policy for continuous-review inventory systems. European Journal of Operational Research 202 (3), 675–685.

Ha, A. Y., 1997. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. Management Science 43 (8), 1093–1103.

Hartman, D. J., 2015. Enhancing and customizing laboratory information systems to improve-enhance pathologist workflow. Surgical Pathology Clinics 8 (2), 137–143.

Higginson, J. K., Bookbinder, J. H., 1994. Policy recommendations for a shipment-consolidation program. Journal of Business Logistics 14 (1), 87–112.

Higginson, J. K., Bookbinder, J. H., 1995. Markovian decision processes in shipment consolidation. Transportation Science 29 (3), 242–255.

Howard, C., 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.

Lieb, R. C., Lieb, K. J., 2015. The North American third-party logistics industry in 2013: The provider CEO perspective. Transportation Journal 54 (1), 104–121.

Lippman, S. A., 1975. Applying a new device in the optimization of exponential queuing systems. Operations Research 23 (4), 687–710.

Merrick, R., Bookbinder, J. H., 2010. Environmental assessment of shipment release policies. International Journal of Physical Distribution and Logistics Management 40 (10), 748–762.

Min, D., 2014. Heuristic procedures for a stochastic batch service problem. 4OR 12 (2), 157–174.

Papadaki, K., Powell, W. B., 2002. Exploiting structure in adaptive dynamic programming algorithms for a stochastic batch service problem. European Journal of Operational Research 142 (1), 108–127.

Papadaki, K., Powell, W. B., 2003. An adaptive dynamic programming algorithm for a stochastic multiproduct batch dispatch problem. Naval Research Logistics 50 (7), 742–769.

Papadaki, K., Powell, W. B., 2007. Monotonicity in multidimensional Markov decision processes for the batch dispatch problem. Operations Research Letters 35 (2), 267–272.

Pazour, J. A., Roy, D., 2015. Analyzing rental vehicle threshold policies that consider expected waiting times for two customer classes. Computers & Industrial Engineering 80, 80–96.

Puterman, M. L., 1994. Markov decision processes: Discrete stochastic dynamic programming. West Sussex, John Wiley and Sons.

Ülkü, M. A., 2012. Dare to care: Shipment consolidation reduces not only costs, but also environmental damage. International Journal of Production Economics 139 (2), 438–446.

Ülkü, M. A., Bookbinder, J. H., 2012. Optimal quoting of delivery time by a third party logistics provider: The impact of shipment consolidation and temporal pricing schemes. European Journal of Operational Research 221 (1), 110–117.

Volmar, K. E., Wilkinson, D. S., Wagar, E. A., Lehman, C. M., 2013. Utilization of stat test priority in the clinical laboratory: a College of American Pathologists q-probes study of 52 institutions. Archives of Pathology & Laboratory Medicine 137 (2), 220–227.

Wilson, R., 2015. Freight moves the economy in 2014. Tech. rep., Council of Supply Chain Management Professionals, CSCMP's Annual State of Logistics Report.

Wong, W. H., Leung, L. C., Hui, Y. V., 2009. Airfreight forwarder shipment planning: A mixed 0-1 model and managerial issues in the integration and consolidation of shipments. European Journal of Operational Research 193 (1), 86–97.

Yang, W., Fung, R. Y., 2014. Stochastic optimization model for order acceptance with multiple demand classes and uncertain demand/supply. Engineering Optimization 46 (6), 824–841.

Yılmaz, Ö., Savaşaneril, S., 2012. Collaboration among small shippers in a transportation market. European Journal of Operational Research 218 (2), 408–415.

Appendix A. Optimality Equation

Let $\pi \in \Pi$ denote a stationary policy for the shipment consolidation and vehicle capacity rationing problem, where Π denotes the set of all feasible stationary policies. Under π, an identical action is taken for the same state in any decision epoch p, i.e., $A_{T_p}^\pi(S) = A^\pi(S) = \pi(S) \ \forall \ p \in \{0, 1, 2, ...\}$ and $S \in \mathbf{S}$. We define $V^\pi(S)$ as the value function under policy π, i.e., the expected discounted total cost when the initial state is S. This value function can be expressed as in Equation (A.1) for a given realization of the initial state at time $T_0 = 0$, i.e., S_0. In the equations below, a) realizations of state variables are denoted by $\mathbf{S_{T_p}}$, and b) the expectations are taken in terms of random variables S_{T_p} and T_p for $p \geq 1$.

$$
\begin{aligned}
V^\pi(\mathbf{S_0}) &= E\Big[\sum_{p=0}^\infty e^{-\alpha T_p} C(S_{T_p}, A^\pi(S_{T_p}))\Big|\mathbf{S_0}\Big] &&\text{(A.1)}\\
&= E\Big[C(\mathbf{S_0}, A^\pi(\mathbf{S_0})) + \sum_{p=1}^\infty e^{-\alpha T_p} C(S_{T_p}, A^\pi(S_{T_p}))\Big|\mathbf{S_0}\Big]\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr} + E\Big[\sum_{p=1}^\infty e^{-\alpha T_p} C(S_{T_p}, A^\pi(S_{T_p}))\Big|\mathbf{S_0}\Big]\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr}\\
&\quad + E\Big[e^{-\alpha T_1} E\Big[\sum_{p=1}^\infty e^{-\alpha(T_p - T_1)} C(S_{T_p}, A^\pi(S_{T_p})\Big|\mathbf{S_0}\Big]\Big] &&\text{(A.2)}\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr}\\
&\quad + E\Big[e^{-\alpha T_1} \sum_{\mathbf{S_{T_1}}\in\mathbf{S}} p(\mathbf{S_{T_1}}|\mathbf{S_0}, A^\pi(\mathbf{S_0})) E\Big[\sum_{p=1}^\infty e^{-\alpha(T_p-T_1)} C(S_{T_p}, A^\pi(S_{T_p}))\Big|\mathbf{S_{T_1}}\Big]\Big]\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr} + E\Big[e^{-\alpha T_1} \sum_{\mathbf{S_{T_1}}\in\mathbf{S}} p(\mathbf{S_{T_1}}|\mathbf{S_0}, A^\pi(\mathbf{S_0}))V^\pi(\mathbf{S_{T_1}})\Big] &&\text{(A.3)}\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr} + E[e^{-\alpha T_1}] \sum_{\mathbf{S_{T_1}}\in\mathbf{S}} p(\mathbf{S_{T_1}}|\mathbf{S_0}, A^\pi(\mathbf{S_0}))V^\pi(\mathbf{S_{T_1}})\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr} + \Big[\int_0^\infty e^{-\alpha t}\lambda e^{-\lambda t}dt\Big]\sum_{\mathbf{S_{T_1}}\in\mathbf{S}} p(\mathbf{S_{T_1}}|\mathbf{S_0}, A^\pi(\mathbf{S_0}))V^\pi(\mathbf{S_{T_1}})\\
&= I_{[A^\pi(\mathbf{S_0})]}K + \frac{1}{\alpha+\lambda}C(\mathbf{S_0} - A^\pi(\mathbf{S_0}))^{tr} + \frac{\lambda}{\alpha+\lambda}\sum_{\mathbf{S_{T_1}}\in\mathbf{S}} p(\mathbf{S_{T_1}}|\mathbf{S_0}, A^\pi(\mathbf{S_0}))V^\pi(\mathbf{S_{T_1}}). &&\text{(A.4)}
\end{aligned}
$$

Equation (A.1) represents the value function under policy π as the expected discounted sum of all one-step costs. Equation (A.2) is obtained by separating the expected cost of the first epoch from the summation in Equation (A.1). Next, the expectation in Equation (A.2) is conditioned first in terms of T_1 (arrival time of the first shipment order) and then in terms of S_{T_1}. Afterwards, Equation (A.3) is obtained by replacing $E\big(\sum_{p=1}^\infty e^{-\alpha(T_p-T_1)} C(S_{T_p}, A^\pi(S_{T_p})|\mathbf{S_{T_1}}\big)$ with $V^\pi(\mathbf{S_{T_1}})$. This replacement is valid

because (i) the series $T_p - T_1, p \geq 1$ and $T_p, p \geq 0$ have the same Erlang distribution for the corresponding terms, (ii) the actions and rewards are stationary. Finally, the expectation of the discount component $(e^{-\alpha T_1})$ in Equation (A.3) is taken to calculate the effective discount factor. We simplify Equation (A.4) by replacing $\mathbf{S_0}$, $A^\pi(\mathbf{S_0})$ and $\mathbf{S_{T_1}}$ with S, $A^\pi(S)$ and S', respectively, and defining $\beta = \lambda/(\alpha + \lambda)$:

$$V^\pi(S) = I_{[A^\pi(S)]}K + \frac{1}{\alpha + \lambda}C(S - A^\pi(S))^{tr} + \beta \sum_{S' \in \mathbf{S}} p(S'|S, A^\pi(S))V^\pi(S'). \quad \text{(A.5)}$$

Equation (A.5) is equivalent to the value function of an infinite-horizon discrete-time total discounted-reward MDP under a stationary policy as described in Equation 6.1.4 of Puterman (1994). Then, the proposed continuous-time model in Equation (3) of the main text can be represented as $V(S) = min_{\pi \in \Pi}V^\pi(S) = V^{\pi^*}(S)$, where π^* is the optimal policy. Therefore, the optimal value function $V(S)$ and the optimal policy π^* can be derived by solving the following recursive Bellman optimality equation for the equivalent infinite-horizon discrete-time discounted Markov decision process (DTMDP).

$$V(S) = min_{A \in \mathbf{A}(S)}\Big\{I_{[A]}K + \frac{1}{\alpha + \lambda}C(S - A)^{tr} + \beta \sum_{S' \in \mathbf{S}} p(S'|S, A)V(S')\Big\}$$

Appendix B. Proofs

Proof of Lemma 1

1) When $k_2 \leq 0$, $A = (a_1, a_2)$ has more Type 1 and Type 2 orders than A'. Let $A - A' = \Delta = (\delta_1, \delta_2) = \delta_1 e_1 + \delta_2 e_2$ s.t., $\delta_i \geq 0$. We define $S' = S + \Delta$. Since $S' \geq S$, both A and A' are feasible actions for S'. Then,

$$V(S, A) = K + \frac{1}{\alpha + \lambda}[c_1(s_1 - a_1) + c_2(s_2 - a_2)] + \beta[\sum_{n=1}^N d_1(n)\frac{\lambda_1}{\lambda}V(s_1 - a_1 + n, s_2 - a_2)$$
$$+ \sum_{m=1}^M d_2(m)\frac{\lambda_2}{\lambda}V(s_1 - a_1, s_2 - a_2 + m)]$$

ii

$$\leq \quad K + \frac{1}{\alpha + \lambda}[c_1([s_1 + \delta_1] - a_1) + c_2([s_2 + \delta_2] - a_2)]$$

$$+\beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V([s_1 + \delta_1] - a_1 + n, [s_2 + \delta_2] - a_2)$$

$$+\sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V([s_1 + \delta_1] - a_1, [s_2 + \delta_2] - a_2 + m)]$$

$$= \quad V(S', A)$$

$$= \quad K + \frac{1}{\alpha + \lambda}[c_1(s_1 - [a_1 - \delta_1]) + c_2(s_2 - [a_2 - \delta_2])]$$

$$+\beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V(s_1 - [a_1 - \delta_1] + n, s_2 - [a_2 - \delta_2])$$

$$+\sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V(s_1 - [a_1 - \delta_1], s_2 - [a_2 - \delta_2] + m)$$

$$= \quad V(S, A')$$

where the first inequality follows by Theorem 1.

2) When $k_2 > 0$,

$$V(S, A) \quad = \quad K + \frac{1}{\alpha + \lambda}[c_1(s_1 - a_1) + c_2(s_2 - a_2)] + \beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V(s_1 - a_1 + n, s_2 - a_2)$$

$$+\sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V(s_1 - a_1, s_2 - a_2 + m)]$$

$$\leq \quad K + \frac{1}{\alpha + \lambda}[c_1(s_1 - a_1 + 1) + c_2(s_2 - a_2 - 1)]$$

$$+\beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V(s_1 - a_1 + n, s_2 - a_2) + \sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V(s_1 - a_1, s_2 - a_2 + m)]$$

iii

$$\leq \quad K + \frac{1}{\alpha + \lambda}[c_1(s_1 - [a_1 - 1]) + c_2(s_2 - [a_2 + 1])]$$
$$+ \beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V([s_1 + 1] - a_1 + n, [s_2 - 1] - a_2)$$
$$+ \sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V([s_1 + 1] - a_1, [s_2 - 1] - a_2 + m)]$$
$$= \quad K + \frac{1}{\alpha + \lambda}[c_1(s_1 - [a_1 - 1]) + c_2(s_2 - [a_2 + 1])]$$
$$+ \beta[\sum_{n=1}^{N} d_1(n)\frac{\lambda_1}{\lambda}V(s_1 - [a_1 - 1] + n, s_2 - [a_2 + 1])$$
$$+ \sum_{m=1}^{M} d_2(m)\frac{\lambda_2}{\lambda}V(s_1 - [a_1 - 1], s_2 - [a_2 + 1] + m)]$$
$$= \quad V(S, A - e_1 + e_2)$$

where the first inequality follows from $c_1 > c_2$, and the second inequality is implied by Theorem 1. By induction, $V(S, A) \leq V(S, A - k_2e_1 + k_2e_2)$ since $0 < k_2 \leq k_1$. By definition and Theorem 1, $V(S, A - k_2e_1 + k_2e_2) \leq V(S + (k_1 - k_2)e_1, A - k_2e_1 + k_2e_2) = V(S, A - k_2e_1 + k_2e_2 - (k_1 - k_2)e_1) = V(S, A - k_1e_1 + k_2e_2)$ since $k_1 \geq k_2$. Therefore $V(S, A) \leq V(S, A - k_1e_1 + k_2e_2)$. □

Proof of Theorem 2

Theorem 1 shows the monotonicity of $V(S)$ in S for any value of ω; therefore, the result is valid for $\omega = \infty$ as well. Similarly, Proposition 1 also applies to the case of $\omega = \infty$. Thus, when $\omega = \infty$, $A^*(S) \in [\overrightarrow{0}, \overline{A}(S)]$, where $\overline{A}(S) = S = (s_1, s_2) \ \forall S \in \mathbf{S}$.

If $A^*(S) = \overline{A}(S)$ for any state $S = (s_1, s_2) \in \mathbf{S}$, the optimal decision is to *SHIP* (i.e., $A^*(S) = S = (s_1, s_2)$). Thus,

$$V(S, \overrightarrow{0}) \geq V(S, \overline{A}(S)) = V(S).$$

In general, the optimal decision in state $S' = S + ie_1 + je_2$ for non-negative i and j is either $\overrightarrow{0}$ or $\overline{A}(S') = S'$ due to Proposition 1. However, the following holds if

$A^*(S) = \overline{A}(S)$.

$$
\begin{aligned}
V(S', \overrightarrow{0}) \;\geq\; & V(S') \\
\geq\; & V(S) = V(S, \overline{A}(S)) \\
=\; & K + \frac{1}{\alpha+\lambda}[c_1(s_1 - s_1) + c_2(s_2 - s_2)] + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha+\lambda} V(s_1 - s_1 + n, s_2 - s_2) \\
& + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha+\lambda} V(s_1 - s_1, s_2 - s_2 + m) \\
\\
=\; & K + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha+\lambda} V(ne_1) + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha+\lambda} V(me_2) \\
=\; & K + \frac{1}{\alpha+\lambda}[c_1(s_1 + i - (s_1 + i)) + c_2(s_2 + j - (s_2 + j))] \\
& + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha+\lambda} V(s_1 + i - (s_1 + i) + n, s_2 + j - (s_2 + j)) \\
& + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha+\lambda} V(s_1 + i - (s_1 + i), s_2 + j - (s_2 + j) + m) \\
=\; & V(S', \overline{A}(S'))
\end{aligned}
$$

The first inequality follows by definition, i.e., $V(S', A) \geq V(S')$ for any $A \in \mathbf{A}(S)$. The second inequality holds due to Theorem 1: $V(S)$ is partially nondecreasing in S; thus, $V(S) \leq V(S')$. Since $V(S', \overrightarrow{0}) \geq V(S', \overline{A}(S'))$, the optimal decision in state S' is $\overline{A}(S')$(i.e., *SHIP*) if $A^*(S) = \overline{A}(S)$. $\square$

Proof of Lemma 2

It is sufficient to prove Lemma 2 for the finite-horizon version of our model since, according to Bertsekas (2001), the result for the finite-horizon model applies to the the infinite-horizon case. Proof is by induction. Consider the finite-horizon version of our model, letting $\mathbf{N}$ be the final decision epoch and $V_{\mathbf{N}}(S) = 0$, $\forall S \in \mathbf{S}$.

We have $V_{\mathbf{N}}(S) = V_{\mathbf{N}}(S + re_1 - qe_2)$, since each $V_{\mathbf{N}}(S)$ is zero for all states.

Step $\mathbf{N} - 1$

$$
\begin{aligned}
V_{\mathbf{N}-1}(S) &= min\{K, \frac{1}{\alpha + \lambda}(c_1 s_1 + c_2 s_2)\}, \quad and \\
V_{\mathbf{N}-1}(S + re_1 - qe_2) &= min\{K, \frac{1}{\alpha + \lambda}(c_1[s_1 + r] + c_2[s_2 - q])\} \\
&= min\{K, \frac{1}{\alpha + \lambda}(c_1 s_1 + c_2 s_2 + rc_1 - \frac{c_1}{q}q)\} \\
&= min\{K, \frac{1}{\alpha + \lambda}(c_1 s_1 + c_2 s_2)\} = V_{\mathbf{N}-1}(S), \quad \forall S \in \mathbf{S}
\end{aligned}
$$

These equations show that for step $\mathbf{N} - 1$, $V_{\mathbf{N}-1}(S) = V_{\mathbf{N}-1}(S + re_1 - qe_2)$.

Step n

Given $V_p(S) = V_p(S + re_1 - qe_2)$ for all $p \in \{n+1, n+2, ..., \mathbf{N}-1, \mathbf{N}\}$ and $S \in \mathbf{S}$, we want to show that $V_n(S) = V_n(S + re_1 - qe_2)$ $\forall S \in \mathbf{S}$. Let $\mathbb{K} = K + \sum_{n'} \frac{d_1(n')\lambda_1}{\alpha + \lambda} V_{n+1}(n'e_1) + \sum_{m'} \frac{d_2(m')\lambda_2}{\alpha + \lambda} V_{n+1}(m'e_2)$. Then:

$$
\begin{aligned}
V_n(S) &= min\{\mathbb{K}, \frac{1}{\alpha + \lambda}(c_1 s_1 + c_2 s_2) + \sum_{n'=1}^{N} \frac{d_1(n')\lambda_1}{\alpha + \lambda} V_{n+1}(S + n'e_1) \\
&\quad + \sum_{m'=1}^{M} \frac{d_2(m')\lambda_2}{\alpha + \lambda} V_{n+1}(S + m'e_2)\} \\
&= min\{\mathbb{K}, \frac{1}{\alpha + \lambda}(c_1 s_1 + c_2 s_2 + [rc_1 - qc_2]) + \sum_{n'=1}^{N} \frac{d_1(n')\lambda_1}{\alpha + \lambda} V_{n+1}(S + n'e_1) \\
&\quad + \sum_{m'=1}^{M} \frac{d_2(m')\lambda_2}{\alpha + \lambda} V_{n+1}(S + m'e_2)\} \\
&= min\{\mathbb{K}, \frac{1}{\alpha + \lambda}(c_1[s_1 + r] + c_2[s_2 - q]) \qquad\qquad\qquad (B.1) \\
&\quad + \sum_{n'=1}^{N} \frac{d_1(n')\lambda_1}{\alpha + \lambda} V_{n+1}(S + n'e_1 + (re_1 - qe_2)) \\
&\quad + \sum_{m'=1}^{M} \frac{d_2(m')\lambda_2}{\alpha + \lambda} V_{n+1}(S + m'e_2 + (re_1 - qe_2))\} \\
&= min\{\mathbb{K}, \frac{1}{\alpha + \lambda}(c_1[s_1 + r] + c_2[s_2 - q]) + \sum_{n'=1}^{N} \frac{d_1(n')\lambda_1}{\alpha + \lambda} V_{n+1}((S + re_1 - qe_2) + n'e_1) \\
&\quad + \sum_{m'=1}^{M} \frac{d_2(m')\lambda_2}{\alpha + \lambda} V_{n+1}((S + re_1 - qe_2) + m'e_2)\} = V_n(S + re_1 - qe_2). \quad \square
\end{aligned}
$$

vi

Equation (B.1) results from the induction condition, i.e., $V_{n+1}(S + ne_1) = V_{n+1}(S + ne_1 + (re_1 - qe_2))$ and $V_{n+1}(S + me_2) = V_{n+1}(S + me_2 + (re_1 - qe_2))$. This proof shows that $V_p(S) = V_p(S + re_1 - qe_2)\ \forall p \in \{1, 2, ..., N-1, N\}$; therefore, Lemma 2 holds as well in the proposed infinite-horizon model based on a result in Bertsekas (2001) (page 8).

Proof of Theorem 3 We need to show that, let $S = (s_1, \overline{s_2}(s_1))$, where $s_1 \in \{0, 1, ...\}$ and $\overline{s_2}(s_1) = min\{s_2 : A^*(s_1, s_2) = \overline{A}(s_1, s_2)\}$. If $\overline{s_2}(s_1) > 0$, then a) $A^*(s_1 + 1, s_2) = \overline{A}(s_1 + 1, s_2),\ \forall s_2\ s.t.\ s_2 \geq \overline{s_2}(s_1) - q$ and b) $A^*(s_1 + 1, s_2) = \overrightarrow{0},\ \forall s_2\ s.t.\ 0 \leq s_2 < \overline{s_2}(s_1) - q$.

Part a) Since $A^*(S) = \overline{A}(S)$ for $S = (s_1, \overline{s_2}(s_1))$, $V(S, \overrightarrow{0}) \geq V(S, \overline{A}(S))$. Thus,

$$
\begin{aligned}
V(S + e_1 - qe_2, \overrightarrow{0}) &= \frac{c_1(s_1 + 1) + c_2(\overline{s_2}(s_1) - q)}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V((S + e_1 - qe_2) + ne_1) \\
&\quad + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V((S + e_1 - qe_2) + me_2) \\
&= \frac{c_1 s_1 + c_2 \overline{s_2}(s_1)}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V((S + e_1 - qe_2) + ne_1) \\
&\quad + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V((S + e_1 - qe_2) + me_2) \\
&= \frac{c_1 s_1 + c_2 \overline{s_2}(s_1)}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V(S + ne_1) \qquad \text{(B.2)} \\
&\quad + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V(S + me_2) \\
&= V(S, \overrightarrow{0}) \\
&\geq V(S, \overline{A}(S)) \\
&= V(S + e_1 - qe_2, \overline{A}(S + e_1 - qe_2))
\end{aligned}
$$

where Equation (B.2) holds due to Lemma 2. This way, we show that $A^*(S + e_1 - qe_2) = \overline{A}(S + e_1 - qe_2)$ since $A^*(S) = \overline{A}(S)$. This result also implies that $A^*(S + e_1 + (k - q)e_2) = \overline{A}(S + e_1 + (k - q)e_2)\ \forall k \in \mathbb{Z}_{\geq 0}$ due to the existence of threshold-type optimal policies (Theorem 2).

Part b) Since $\overline{s_2}(s_1) = min\{s_2 : A^*(S) = \overline{A}(S)\}$, $V(S - e_2, \overrightarrow{0}) \leq V(S - e_2, \overline{A}(S - e_2))$.

Then,

$$
\begin{aligned}
V(S + e_1 - (q+1)e_2, \overrightarrow{0}) &= \frac{c_1(s_1 + 1) + c_2(\overline{s_2}(s_1) - (q+1))}{\alpha + \lambda} \\
&\quad + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V((S + e_1 - (q+1)e_2) + ne_1) \\
&\quad + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V((S + e_1 - (q+1)e_2) + me_2) \\
\\
&= \frac{c_1 s_1 + c_2(\overline{s_2}(s_1) - 1)}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V((S - e_2 + ne_1) + e_1 - qe_2) \\
&\quad + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V((S - e_2 + ne_1) + e_1 - qe_2) \\
&= \frac{c_1 s_1 + c_2(\overline{s_2}(s_1) - 1)}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V(S - e_2 + ne_1) \\
&\quad + \sum_{m=1} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V(S - e_2 + me_2) \qquad\qquad (B.3) \\
&= V(S - e_2, \overrightarrow{0}) \\
&\leq V(S - e_2, \overline{A}(S - e_2)) \\
&= V((S - e_2) + e_1 - qe_2, \overline{A}((S - e_2) + e_1 - qe_2)).
\end{aligned}
$$

Equation (B.3) holds due to Lemma 2: $V((S - e_2 + ne_1) + e_1 - qe_2) = V(S - e_2 + ne_1)$ and $V((S - e_2 + me_2) + e_1 - qe_2) = V(S - e_2 + me_2)$. This shows that $A^*(S - e_1) = \overrightarrow{0}$ implies $V(S + e_1 - (q+1)e_2, \overrightarrow{0}) \leq V((S - e_2) + e_1 - qe_2, \overline{A}((S - e_2) + e_1 - qe_2))$. Therefore, $A^*((S - e_1) + e_1 - qe_2) = \overrightarrow{0}$, which guarantees that $A^*(S + e_1 - (q+k)e_2) = \overrightarrow{0}$, $\forall k, 0 \leq k \leq \overline{s_2}(s_1) - q$ by Theorem 2. $\square$

Proof of Lemma 3

Let S be a feasible state with $s_1 c_1 + s_2 c_2 \geq K(\alpha + \lambda)$. Then:

$$
\begin{aligned}
V(S, \overrightarrow{0}) &= \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V(S + ne_1) + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V(S + me_2) \\
&\geq K + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V(S + ne_1) + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V(S + me_2) \\
&\geq K + \sum_{n=1}^{N} \frac{d_1(n)\lambda_1}{\alpha + \lambda} V(ne_1) + \sum_{m=1}^{M} \frac{d_2(m)\lambda_2}{\alpha + \lambda} V(me_2) \qquad \text{(B.4)} \\
&= V(S, \overline{A}(S))
\end{aligned}
$$

where Equation (B.4) follows by Theorem 1. This implies that $A^*(S) = \overline{A}(S)$ for any $S = (0, s_2)$, where $s_2 \geq s_2^{UB}$, because $c_2 s_2 \geq c_2 s_2^{UB} = K(\alpha + \lambda)$. Therefore, $\overline{s_2}(0) = min\{s_2 : A^*(0, s_2) = \overline{A}(0, s_2)\}$ has to be less than or equal to s_2^{UB}. $\square$

Proof of Lemma 4

When S is small enough (e.g. $(0, 1)$, $(1, 0)$), K could be significantly larger than $\frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda}$, so that the optimal decision could be to *WAIT*. In such a case, $\frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} V(S + e_1) + \frac{\lambda_2}{\alpha + \lambda} V(S + e_2) \leq K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2)$ holds under the assumption of $d_1(1) = d_2(1) = 1$. Then:

$$
\begin{aligned}
V(S) &= \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} V(S + e_1) + \frac{\lambda_2}{\alpha + \lambda} V(S + e_2) \\
&\geq \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} V(S + e_2) + \frac{\lambda_2}{\alpha + \lambda} V(S + e_2) \qquad \text{(B.5)} \\
&= \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} V(S + e_2) \\
&\geq \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} V(S) \Rightarrow \qquad \text{(B.6)} \\
V(S) &\geq \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} V(S) \qquad \text{(B.7)}
\end{aligned}
$$

Inequality (B.5) holds because V(S) is monotone non-decreasing w.r.t. shipment order type (i.e., $V(S + e_1) \geq V(S + e_2)$ since $S + e_1 \succeq S + e_2$ as indicated in Theorem 1). Inequality (B.6) follows from Theorem 1 ($V(S + e_2) \geq V(S)$ as $S + e_2 > S$). By rearranging Inequality (B.7), we obtain that $V(S) \geq \frac{c_1 s_1 + c_2 s_2}{\alpha}$ when $A^*(S) = \overrightarrow{0}$. Therefore, when $A^*(e_1) = A^*(e_2) = \overrightarrow{0}$, $V(e_i) \geq \frac{c_i}{\alpha}$ for $i \in \{1, 2\}$.

Let $(0, \overline{s_2}(0))$ be at the border of the optimal linear-staircase threshold for shipment

decisions, i.e., $A^*(S) = \overline{A}(S)$ and $A^*(S-1) = \overrightarrow{0}$, for $S = (0, \overline{s_2}(0))$. Then:

$$
\begin{aligned}
V(0, \overline{s_2}(0)) &= K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \\
&\leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} V(S + e_1) + \frac{\lambda_2}{\alpha + \lambda} V(S + e_2) \qquad (\text{B}.8) \\
&= \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} \left[K + \frac{\lambda_1}{\alpha + \lambda}) V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \right] \Rightarrow \qquad (\text{B}.9)
\end{aligned}
$$

$$
K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} \left[K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \right] \Rightarrow
$$

$$
0 \leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} - (1 - \frac{\lambda}{\alpha + \lambda}) \left[K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \right]. \qquad (\text{B}.10)
$$

Inequality (B.8) holds because the optimal action is to $SHIP$ in state S (i.e., $A^*(S) = \overline{A}(S)$). Equation (B.9) follows from Theorem 2, which implies that $A^*(S + e_i) = \overline{A}(S)$ for $i \in \{1, 2\}$ because $S + e_i > S$ and $A^*(S) = \overline{A}(S)$. Thus, $V(S + e_1) = V(S + e_1) = K + \left(\frac{\lambda_1}{\alpha + \lambda} \right) V(e_1) + \left(\frac{\lambda_2}{\alpha + \lambda} \right)$. We derive the upper bound from Inequality (B.10) as follows:

$$
\begin{aligned}
\overline{s_2}(0) &\geq \frac{\alpha + \lambda}{c_2} \left(1 - \frac{\lambda}{\alpha + \lambda} \right) \left[K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \right] \\
&= \frac{\alpha}{c_2} \left[K + \frac{\lambda_1}{\alpha + \lambda} V(e_1) + \frac{\lambda_2}{\alpha + \lambda} V(e_2) \right] \\
&\geq \frac{\alpha}{c_2} \left[K + \frac{\lambda_1}{\alpha + \lambda} \frac{c_1}{\alpha} + \frac{\lambda_2}{\alpha + \lambda} \frac{c_2}{\alpha} \right] \\
&= \frac{1}{c_2} \left[\alpha K + \frac{c_1 \lambda_1}{\alpha + \lambda} + \frac{c_2 \lambda_2}{\alpha + \lambda} \right]
\end{aligned}
$$

Therefore $\frac{1}{c_2} \left[\alpha K + \frac{c_1 \lambda_1}{\alpha + \lambda} + \frac{c_2 \lambda_2}{\alpha + \lambda} \right]$ gives a lower bound on $\overline{s_2}(0)$. $\square$

Appendix C. Extending Lemma 5 for Batch Arrivals

Let $(0, \overline{s_2}(0))$ be at the border of the optimal linear-staircase threshold for shipment decisions. When $\omega = \infty$ and $A^*(Ne_1) = A^*(Me_2) = \overrightarrow{0}$, $\overline{s_2}(0) \geq s_2^{LB} = \frac{1}{c_2} \left[\alpha K + \frac{c_1 \lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} n d_1(n) + \frac{c_2 \lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} m d_2(m) \right]$.

Proof: This proof follows the same logic as that of Lemma 4. Suppose that $V(S) =$

$V(S, \overrightarrow{0})$ for some state S. Then for those states, the following holds:

$$
\begin{aligned}
V(S) &= \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(S + ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(S + me_2) \\
&\geq \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(S) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(S) \\
&= \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} V(S) \Rightarrow \\
V(S) &\geq \frac{c_1 s_1 + c_2 s_2}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} V(S) \Rightarrow \\
V(S) &\geq \frac{c_1 s_1 + c_2 s_2}{\alpha}
\end{aligned}
\tag{C.1}
$$

Since $A^*(Ne_1) = A^*(Me_2) = \overrightarrow{0}$, $A^*(ne_1) = \overrightarrow{0}$ and $A^*(me_2) = \overrightarrow{0}$ $\forall n \in \{1, \dots N\}, m \in \{1, \dots M\}$ due to Theorem 3. Therefore, $V(ne_1) \geq \frac{nc_1}{\alpha}$ and $V(me_2) \geq \frac{mc_2}{\alpha}$ $\forall n \in \{1, \dots N\}, m \in \{1, \dots M\}$. In addition, $A^*(\overline{s_2}(0)e_2) = \overline{A}(\overline{s_2}(0)e_2)$ and $A^*((\overline{s_2}(0) - 1)e_2) = \overrightarrow{0}$ because $(0, \overline{s_2}(0))$ is at the border of the optimal shipment threshold . Then:

$$
\begin{aligned}
V(\overline{s_2}(0)e_2) &= V(\overline{s_2}(0)e_2, \overline{A}(\overline{s_2}(0)e_2)) = K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(me_2) \\
&\leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(\overline{s_2}(0)e_2 + ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(\overline{s_2}(0)e_2 + me_2) \\
&= \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) \left[K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n'=1}^{N} d_1(n') V(n'e_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m'=1}^{M} d_2(m') V(m'e_2) \right] \\
&\quad + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(n) \left[K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n'=1}^{N} d_1(n') V(n'e_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m'=1}^{M} d_2(m') V(m'e_2) \right] \\
&= \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \frac{\lambda}{\alpha + \lambda} \left[K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n'=1}^{N} d_1(n') V(n'e_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m'=1}^{M} d_2(m') V(m'e_2) \right] \Rightarrow \\
K &+ \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(me_2) \\
&\leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} + \left(\frac{\lambda}{\alpha + \lambda} \right) \left[K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(me_2) \right] \Rightarrow \\
0 &\leq \frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} - \left(1 - \frac{\lambda}{\alpha + \lambda} \right) \left[K + \frac{\lambda_1}{\alpha + \lambda} \sum_{n=1}^{N} d_1(n) V(ne_1) + \frac{\lambda_2}{\alpha + \lambda} \sum_{m=1}^{M} d_2(m) V(me_2) \right] \Rightarrow
\end{aligned}
\tag{C.2}
$$

xi

$$\frac{c_2 \overline{s_2}(0)}{\alpha + \lambda} \geq \frac{\alpha}{\alpha + \lambda}\left[K + \frac{\lambda_1}{\alpha + \lambda}\sum_{n=1}^{N} d_1(n)V(ne_1) + \frac{\lambda_2}{\alpha + \lambda}\sum_{m=1}^{M} d_2(m)V(me_2)\right] \Rightarrow$$

$$\overline{s_2}(0) \geq \frac{\alpha}{c_2}\left[K + \frac{\lambda_1}{\alpha + \lambda}\sum_{n=1}^{N} d_1(n)V(ne_1) + \frac{\lambda_2}{\alpha + \lambda}\sum_{m=1}^{M} d_2(m)V(me_2)\right]$$

$$\geq \frac{\alpha}{c_2}\left[K + \frac{\lambda_1}{\alpha + \lambda}\sum_{n=1}^{N} d_1(n)\frac{nc_1}{\alpha} + \frac{\lambda_2}{\alpha + \lambda}\sum_{m=1}^{M} d_2(m)\frac{mc_2}{\alpha}\right]$$

$$= \frac{1}{c_2}\left[\alpha K + \frac{c_1\lambda_1}{\alpha + \lambda}\sum_{n=1}^{N} nd_1(n) + \frac{c_2\lambda_2}{\alpha + \lambda}\sum_{m=1}^{M} md_2(m)\right]$$

$$= \frac{1}{c_2}\left[\alpha K + \frac{c_1\lambda_1}{\alpha + \lambda}D_1 + \frac{c_2\lambda_2}{\alpha + \lambda}D_2\right] \ \square$$

Appendix D. Customized Solution Approach

The customized solution approach is proposed for the case of $\omega = \infty$, $c_1 = qc_2$, and batch arrivals, considering the equivalent discrete-time infinite-horizon MDP formulation in Equation (6). In this case, the optimal decisions are of control limit type, with an optimal linear-staircase threshold separating the states (combinations of s_1 and s_2 quantities) for $WAIT$ and $SHIP$ decisions based on Theorem 3. This linear-staircase threshold has a step length of q. Under the optimal linear-staircase threshold policy, the states for $WAIT$ decisions establish a **triangle**-like $WAIT$ region as shown in Figure 2, while the remaining states form a $SHIP$ region. Therefore, when the system state enters the $SHIP$ region from the $WAIT$ region after a series of state transitions, the optimal policy makes a shipment (dispatches the consolidated load on a vehicle), which changes the system state to $(0,0)$ until the next shipment order arrival.

Under any linear-staircase threshold policy, the system executes a *cycle* by starting from state $(0,0)$ and returning back to this state after a random duration, due to a shipment caused by leaving the $WAIT$ region. This cycle is repeated indefinitely; thus, our model behaves like a renewal process under a linear-staircase threshold policy. Due to this renewal-process-like behaviour, we can derive a formulation for $V^{\pi(s_2)}(0,0)$, the expected total discounted cost under the linear-staircase threshold policy that starts at state $(0, s_2)$ with step length q given that no orders await shipment initially. $V^{\pi(s_2)}(0,0)$ depends on a) the expected cost of the cycle described above, and 2) the distribution of the number of decision epochs within the cycle. The *optimal* linear-staircase threshold policy can be determined by finding the s_2 value in the range of $[s_2^{LB}, s_2^{UB}]$ that minimizes $V^{\pi(s_2)}(0,0)$. The s_2 value derived from this optimization is equal to $\overline{s_2}(0)$.

Now let us derive a formulation for $V^{\pi(s_2)}(0,0)$ under the linear-staircase threshold policy starting at state $(0, s_2)$. Note that all variables defined below are specific to this

threshold policy. Although most of the components in the formulation below depends on s_2 choice, we omit it from the notation for simplicity. Therefore, one needs to define all the components below anew, when he/she uses the $V^{\pi(s_2)}(0,0)$ formulation for each s_2 in the range of $[s_2^{LB}, s_2^{UB}]$.

We define $\mathbf{F}$ as the set of states in which a cycle may end under the considered linear-staircase threshold policy. Any state (i,j) in the $SHIP$ region may be an element of $\mathbf{F}$ if $(i-N, j)$ or $(i, j-M)$ is in the $WAIT$ region formed by the linear-staircase threshold policy, where $N = max(n : d_1(n) > 0)$ and $M = max(m : d_2(m) > 0)$. For instance, $\mathbf{F} = \{(0,17), (1,16), (1,15), (2,14), (2,13) \ldots (8,2), (8,1), (9,0)\}$ for the example illustrated in Figure 2. We aim to calculate the expected discounted cost of the cycle for each state in set $\mathbf{F}$.

Let us consider a cycle starting in state $(0,0)$ and ending in state $(i,j) \in \mathbf{F}$. This cycle must visit a state (i', j') in the $WAIT$ region at the decision epoch before the last one (in the cycle). The system may follow different state transition pathways to reach state (i', j') from state $(0,0)$. Each of those pathways may result in a different holding cost, depending on the arrival sequence and sizes of shipment orders, which are tracked by X_n and Y_m (the total number of Type 1 arrivals of order-size n and Type 2 arrivals of order-size m until reaching state (i', j'), respectively). We define (X, Y) as the vector of X_n and Y_m variables for $n \in \{1, ...N\}$ and $m \in \{1, ...M\}$, i.e., $(X, Y) = \big((X_1, X_2, ...X_N), (Y_1, Y_2, ...Y_M)\big)$ for the destination state $(i', j') \in \mathbf{F}$ where $i' = \sum_{n=1}^{N} n X_n$ and $j' = \sum_{m=1}^{M} m Y_m$. Then, $\Psi(i', j') = \sum_{n=1}^{N} X_n + \sum_{m=1}^{M} Y_m$ is the total number of decision epochs until reaching state (i', j') for the group of state transition pathways represented by vector (X, Y). We also define $L(i', j')$ as the set of all feasible (X, Y) vectors for state (i', j').

The most critical step in calculating the total discounted cost of a particular linear-staircase threshold policy is specifying the holding cost and the associated discount factor for each decision epoch. For example, the last order arrival that triggers a shipment does not incur any holding cost. If the shipment order that arrives before the last one is Type 1 with size n, then those n orders are held in the system for one epoch and incur a cost of $\frac{nc_1}{\alpha + \lambda} \beta^{\Psi_{X,Y}}$. If the $(\Psi_{X,Y} - 1)^{th}$ order arrival is Type 2 with size m, then those m orders are held in the system for two epochs incurring a cost of $\frac{mc_2}{\alpha + \lambda}(\beta^{\Psi_{X,Y} - 1} + \beta^{\Psi_{X,Y}})$ until the end of the cycle. In short, the effective discount rate for the holding cost of the $(\Psi_{X,Y} - k)^{th}$ order arrival will be $\beta^{\Psi_{X,Y} - k} + \ldots + \beta^{\Psi_{X,Y}}$.

The vector $(X, Y) \in L(i', j')$ represents $\binom{\Psi_{X,Y}}{X_1, X_2, ..., X_N, Y_1, Y_2, ..., Y_M}$ different state transition pathways starting from state $(0,0)$ and ending up in state (i', j'). Among them, $\binom{\Psi_{X,Y} - 1}{X_1 - 1, X_2, ..., X_N, Y_1, Y_2, ..., Y_M}$ pathways have a shipment order arrival of size one at a partic-

ular decision epoch. By dividing the latter number into the former one, the probability of having an order of size one at the k^{th} decision epoch in the cycle is calculated as $X_1/\Psi_{X,Y}$ for $k \in \{1, 2, ... \Psi_{X,Y}\}$. Following the same logic, $X_n/\Psi_{X,Y}$ and $Y_m/\Psi_{X,Y}$ are the probabilities of observing an order arrival of size n and m at any given decision epoch for vector $(X, Y) \in \mathrm{L}(i', j')$. We denote the expected total discounted holding cost incurred until reaching state (i', j') from state $(0,0)$ for a given $(X, Y) \in \mathrm{L}(i', j')$ as $g_\beta(X, Y)$. This can be derived as follows:

$$g_\beta(X,Y) = \Big[\sum_{n=1}^{N}\sum_{l=1}^{\Psi_{X,Y}}\sum_{u=l}^{\Psi_{X,Y}} nX_n\frac{c_1}{\alpha+\lambda}\beta^u + \sum_{m=1}^{M}\sum_{l=1}^{\Psi_{X,Y}}\sum_{u=l}^{\Psi_{X,Y}} mY_m\frac{c_2}{\alpha+\lambda}\beta^u\Big]/\Psi_{X,Y}$$

The probability that the system state follows the transition pathways represented by vector $(X, Y) \in \mathrm{L}(i', j')$ to reach state (i', j') is given by:

$$P(X,Y) = \Pi_{n=1}^{N}\left[\left(\frac{d_1(n)\lambda_1}{\lambda}\right)^{X_n}\right] \times \Pi_{m=1}^{M}\left[\left(\frac{d_2(m)\lambda_2}{\lambda}\right)^{Y_m}\right] \times \begin{pmatrix} \Psi_{X,Y} \\ X_1, X_2, ..., X_N, Y_1, Y_2, ..., Y_M \end{pmatrix}$$

Now, we are ready to calculate $V^{\pi(s_2)}(0,0)$ by simultaneously specifying at which state (i, j) the cycle ends, what is the previous state in the wait region, and which group of state transitions (represented by (X,Y)) are followed. In the formulation of $V^{\pi(s_2)}(0,0)$ given below, $IW(i', j')$ is an indicator function which is equal to one if state (i', j') is in the $WAIT$ region formed by the considered linear-staircase threshold policy.

$$V^{\pi(s_2)}(0,0) =$$
$$\sum_{(i,j)\in\mathbf{F}}\sum_{n=1}^{N}\sum_{(X,Y)\in\mathrm{L}(i-n,j)} IW(i-n,j)\frac{d_1(n)\lambda_1}{\lambda}P(X,Y)\big[g_\beta(X,Y)+\beta^{1+\Psi_{X,Y}}(K+V^{\pi(s_2)}(0,0))\big]$$
$$+ \sum_{(i,j)\in\mathbf{F}}\sum_{m=1}^{M}\sum_{(X,Y)\in\mathrm{L}(i,j-m)} IW(i,j-m)\frac{d_2(m)\lambda_2}{\lambda}P(X,Y)\big[g_\beta(X,Y)+\beta^{1+\Psi_{X,Y}}(K+V^{\pi(s_2)}(0,0))\big] \Rightarrow$$
$$V^{\pi(s_2)}(0,0) =$$
$$\Big(\sum_{(i,j)\in\mathbf{F}}\sum_{n=1}^{N}\sum_{(X,Y)\in\mathrm{L}(i-n,j)} IW(i-n,j)\frac{d_1(n)\lambda_1}{\lambda}P(X,Y)\big[g_\beta(X,Y)+\beta^{1+\Psi_{X,Y}}K\big]$$
$$+ \sum_{(i,j)\in\mathbf{F}}\sum_{m=1}^{M}\sum_{(X,Y)\in\mathrm{L}(i,j-m)} IW(i,j-m)\frac{d_2(m)\lambda_2}{\lambda}P(X,Y)\big[g_\beta(X,Y)+\beta^{1+\Psi_{X,Y}}K\big]\Big)/h_\beta,$$

$$\text{where } h_\beta = 1 - \Big(\sum_{(i,j)\in\mathbf{F}}\sum_{n=1}^{N}\sum_{(X,Y)\in\mathrm{L}(i-n,j)}\big[\frac{d_1(n)\lambda_1}{\lambda}IW(i-n,j)P(X,Y)\beta^{1+\Psi_{X,Y}}\big]$$
$$+ \sum_{(i,j)\in\mathbf{F}}\sum_{m=1}^{M}\sum_{(X,Y)\in\mathrm{L}(i,j-m)}\big[\frac{d_2(m)\lambda_2}{\lambda}IW(i,j-m)P(X,Y)\beta^{1+\Psi_{X,Y}}\big]\Big)$$

xiv

Note that the complexity of the above $V^{\pi(s_2)}(0,0)$ formulation lies in defining all feasible (X,Y) vectors for the destination state (i,j). The number of feasible (X,Y) vectors could be large for large values of ω, N, and M. However, $V^{\pi(s_2)}(0,0)$ is very easy to calculate when shipment orders are composed of standard-size packages, e.g., $d_1(1)=d_2(1)=1$. That is because, in this case, $L(i',j')$ contains a single (X,Y) vector which is $(X_1=i', Y_1=j')$.

Appendix E. Search Protocol for Alternative Quantity Policies

The proposed MDP can be used to identify such policies in cases similar to that of UPS-Turkey. For this purpose, one should adjust the holding costs in a systematic way and force the MDP model to choose an alternative policy with a less aggressive shipment schedule. The true performance of the alternative policies should be measured using the simulation model described in Section 4.2 under the original c_1 & c_2 values. We derived such an alternative policy for the case of UPS-Turkey using the following search logic:

Step 0: Specify a reasonable step size, z.

Step 1: Keep ρ constant, and reduce c_1 by z. Then, run the MDP model to find a potential alternative policy.

Step 2: Evaluate the performance of the policy from Step 1 using the simulation model described in Section 4.2 under the original holding costs. If both total and shipment costs are reduced compared to the benchmark time-policy stop; go to Step 3. Otherwise go back to Step 1.

Step 3. If the policy derived in Step 2 needs to have a better rate of timely-shipped orders, increase only c_1 by z (let ρ vary, keep c_2 constant), and run the MDP model.

Step 4. Evaluate the performance of the policies from Step 3. If the rate of timely shipped policies is satisfactory and both costs are lower, stop. Otherwise go to Step 3.

The search mechanism in Step 1 (keeping ρ constant and reducing c_1 & c_2) increases the importance of shipment cost in the objective function, and thus, leads to a quantity policy with better shipment cost reduction. The search mechanism in Step 3 (keep c_2 constant and increase c_1) increases the importance of timely-shipment of expedited orders, thus improves the due date compliance. Note that value of step size z depends on the problem type; thus, an appropriate z value should be specified via initial testing. Following the search protocol described above, we found an alternative quantitative policy (with around 20% reduced c_1 & c_2) that achieves the desired performance criteria for the case of UPS-Turkey. The performance of the alternative quantity policy under the original holding cost setting is presented in Table E.2.

Table E.2: % Performance differences of the Alternative and Original Optimal Policies Compared to the Time Policy under the Original Holding Costs.

	% Timely-shipped Orders	Average Lateness (Late Orders)	Average Lateness
Alternative Quantity Policy			
Expedited Orders	0.86%	1.8 %	-13.16%
Standard Orders	0.2%	26.19 %	25%

	Shipment Cost	Holding Cost	Total Cost
% Cost Improvements over Time Policy			
Alternative Quantity Policy	-2.18%	-3.23%	-2.75%
Original Optimal Quantity Policy	6.26%	-10.95%	-3.20%

xvi