

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Carrier Portfolio Management for Shipping Seasonal Products

Tao Lu

Rotterdam School of Management, Erasmus University, Rotterdam, The Netherlands, lu@rsm.nl

Jan C. Fransoo

School of Industrial Engineering, Eindhoven University of Technology, The Netherlands j.c.fransoo@tue.nl

Chung-Yee Lee

Department of Industrial Engineering and Logistics Management, Hong Kong University of Science and Technology, Hong Kong, cylee@ust.hk

Many seasonal products are transported via ocean carriers from origin to destination markets. The shipments arriving earlier in the market may sell at higher prices, but faster shipping services can be costly. In this paper, we study a newsvendor-type shipper who transports and sells seasonal products to an overseas market, where the selling price declines over time. A set of vessels with different schedules and freight rates are available to choose from. Our analysis demonstrates that a portfolio of vessels has two distinct effects on mitigating uncertainties in both demand and vessels' arrival schedules, while these two portfolio effects have been previously understood as separate issues in the literature. To find the optimal portfolio in our problem, we first show that when vessels arrive in a deterministic sequence, the optimal portfolio can either be derived in closed form (in the single-demand setting) or computed efficiently with a variation of the shortest-path algorithm (in the multi-demand setting). Then, based on these results, we propose an approximation procedure to address the general problem with an uncertain arrival sequence. In each iteration of the procedure, we only need to minimize a cost function approximated by a deterministic arrival schedule and the portfolio generated can converge to the optimal one under mild conditions. Finally, we present a real-world case study to demonstrate several practical implications of managing a carrier portfolio.

Key words: ocean transport; diversification; carrier portfolio; newsvendor problem.

1. Introduction

Exporters of seasonal products often face the dilemma of choosing from various ocean shipping services. In reality, shipping services between the same set of origin and destination may differ significantly in the transit time and freight rate due to carriers' different network structures and steaming strategies. For instance, Van de Weijer (2013) reports that among shipping services from

Manila to Rotterdam in 2011, the fastest service required an average transit time of 28.4 days, whereas the average transit time of the slowest service was more than 40 days. Impressively, the slow service costs only half as much as the fast one.

Delivery speed is clearly an important consideration, for example, for fruit exporters. Zespri International Ltd., a New Zealand exporter, delivers fruit to overseas markets via ocean container transport. Their transport-division chief once commented that “getting volumes of the fruit to their destination within the first two weeks of picking is paramount to Zespri ensuring and maintaining its global market share” (Mechelen 2011). On the one hand, the wholesale price in the fruit market depends on a fruit’s degree of freshness/deterioration (Agricultural Marketing Service 2012). On the other hand, given a fruit’s short lifetime, late delivery will leave downstream retailers less time to make sales. Thus, fruit exporters must often compensate their customers for the opportunity cost resulting from late arrivals. Hence, the selling season from the exporter’s perspective is even shorter than the actual lifetime of a fruit. Moreover, a fruit usually has a fixed harvest time. Once the kiwifruit in New Zealand ripens in early May each year, they must be picked as soon as possible, then stored and sent to the distributors in destination markets (e.g., Europe). This requirement makes the speed-cost tradeoff especially critical for fruit exporters like Zespri: While chartered vessels can sail faster than ocean liners and enable the company to sell the fruit at a premium price in the early season, the transportation cost would be substantial and may outweigh the gains. Recently, Zespri has started using a combination of chartered vessels and slow but cheaper ocean liners (Zespri 2014, p. 16), even though shipments in the latter may arrive late.

The risks confronting fruit exporters are twofold. First, due to the long lead time for ocean transport, fruit exporters must make shipping decisions long before demand uncertainty is resolved. For instance, the fruit from New Zealand, after being loaded on board, would spend 4 to 6 weeks at sea before arriving at the Port of Rotterdam. Second, ocean shipping services are notoriously unreliable. For instance, in 2014, severe congestion at the Port of Manila resulted in an average vessel delay of more than five days (Lavigne 2014). Given these uncertainties, it remains unclear how exporters should select from potential vessels with different freight rates and (departure and arrival) schedules, in order to balance the trade-off between speed and cost. What are the benefits for a shipper of using a portfolio of shipping services?

This paper will address the above issues and uncover the potential benefits of using a carrier portfolio. We develop a newsvendor-type model in which a shipper delivers a type of seasonal product to an overseas destination market via ocean carriers. The market value of the product declines over time. Given a set of potential vessels with different schedules and freight rates, the shipper decides which vessel(s) to use and the shipping volume to be assigned to each of the selected vessels in the presence of both demand and arrival schedule uncertainties.

By explicitly solving a special case with two candidate vessels and a single demand that can be filled once shipments arrive, we show that a portfolio of shipping services has two distinct effects on mitigating uncertainties in *both* demand and arrival times. First, a combination of fast and slow services can be used to manage demand uncertainty. The fast service brings a high profit margin, whereas the slow one contributes to a high service level throughout the entire season. Interestingly, each shipping service may be viewed as a supply option contract studied in Martínez-de Albéniz and Simchi-Levi (2005, 2009) with a reservation cost and a *random* execution profit arising from the uncertain arrival schedule. Second, when vessels arrive in an uncertain sequence, splitting shipments between vessels can also help spread the risk from unreliable service schedules. This diversification effect mirrors that studied in the supply risk literature (e.g. Anupindi and Akella 1993, Dada et al. 2007) in which multiple suppliers are sourced to mitigate uncertain supply. In some sense, our shipping procurement provides a unified view of newsvendor problems with supply option contracts and supply diversification. While the above two portfolio effects are treated as separate issues in the inventory literature, in the shipping practice they occur concurrently. Yet, to our knowledge, there are no existing models/applications in which these two effects are combined.

Besides the two portfolio effects, with multiple demands occurring over time, using vessels with late departures may also help manage the inventory costs incurred at the origin and the destination. Our shipping problem combines all these features. Little is known about how to determine the optimal procurement strategy when all these portfolio effects coexist. We start our analysis by assuming that vessels arrive in a deterministic sequence. Under this assumption, the optimal portfolio can either be derived in a closed form with an efficient-frontier characterization (in the single-demand setting), or computed in polynomial time with a variation of shortest-path algorithms (in the multi-demand setting). Based on these results, we propose an iterative approximation procedure to solve the general model with an uncertain arrival sequence. Our solution approach, partly built on several seminal papers such as Martínez-de Albéniz and Simchi-Levi (2005), Fu et al. (2010) and Cheung and Powell (2000), is able to asymptotically generate the optimal portfolio under mild conditions. Finally, we present a numerical study based on the Zespri case.

Our contributions include (1) developing a newsvendor model motivated by a new application in the ocean shipping industry, (2) demonstrating that our application combines two portfolio effects which have been previously known as separate issues, (3) analyzing and solving the model by extending existing results in the literature, and (4) conducting a real-world case study to demonstrate the practical implications of our model.

The rest of this paper is organized as follows. We review related literature in Section 2, and analyze a special case with only two shipping services in Section 4 in order to demonstrate the two effects of a carrier portfolio. In Section 5, we propose solution approaches to the general problem.

Section 6 presents the case study and Section 7 discusses several extensions. Finally, Section 8 concludes.

2. Literature Review

In the literature on maritime logistics, Wan and Levary (1995) develop a method for the shipper to negotiate the freight rate with the shipping companies. Lim et al. (2008) study the shipper's procurement problem in which the shipper faces seasonally varying demands and must comply with the volume guarantees given in the shipping contract. In their model, the shipper manages a network of shipping lanes, but the demand on each lane is known. For inland transportation, Van de Weijer (2013) studies an inbound supply chain of a distribution center in the hinterland. Based on an inventory model, they simulate the effects of hinterland modal split on the transport and inventory costs. Sharypova et al. (2012) develop an analytical model to address the coordination problem between the shipper and inland terminals. They identify a trade-off between variable transportation costs and inventory holding cost in the hinterland supply chain system.

As mentioned earlier, our shipping problem is related to two streams of literature on newsvendor procurement. First, the portfolio effect with respect to demand uncertainty mirrors that studied in Martínez-de Albéniz and Simchi-Levi (2005, 2009). The relations between our model and theirs will be detailed in Section 4. Second, the effect on mitigating uncertain arrival schedules is related to the large body of literature on supply diversification (e.g., Anupindi and Akella 1993, Tomlin 2006, Dada et al. 2007, Federgruen and Yang 2008, 2009). Also see Tomlin and Wang (2010) and Snyder et al. (2010) for comprehensive surveys. Most of the papers along this stream are concerned with the supply uncertainty that leads to a shortage in the supply volume. In our shipping problem, however, delayed shipments will eventually be received but will have to be sold at a lower price, depending on the length of delay. Anupindi and Akella (1993) consider a multi-period problem where the undelivered quantity in the current period will be delivered one period later in their Model III. Our model differs from theirs in that ours captures the trade-off between transit times and freight rates, a major challenge faced by shippers in ocean transport.

Another stream of literature studies multi-period inventory control when multiple delivery modes with deterministic lead times are available. Fukuda (1964) derives the optimal policies when the lead times of two delivery modes differ by exactly one period. Whittemore and Saunders (1977) consider the case where the lead times of two modes vary by an arbitrary number of periods. Feng et al. (2006) further show that with multiple delivery modes base-stock policies are not optimal. Motivated by the transport of short-lifetime products, we depart from the above papers by focusing on a single selling season which is divided into periods with markdowns. Moreover, uncertain arrival times are considered in our model.

For stochastic lead times, pioneering papers like Kaplan (1970) and Zipkin (1986) focus on characterizing optimal policies under sole sourcing. Ramasesh et al. (1991, 1993) and Kouvelis and Li (2008) study dual sourcing strategies with uncertain lead times when demand arrives at a constant rate. Under both demand and lead time uncertainty, existing papers rely mostly on numerical simulations (e.g. Lau and Zhao 1993, Sedarage et al. 1999).

3. The Model

We consider a shipper selling a single type of seasonal product (e.g., fruit) to overseas customers (e.g., distributors and wholesalers in the destination wholesale market). The product is transported from the origin to the destination market via ocean carriers.

The production completes and the product is ready to ship at a given time point, which will be referred to as the “harvest time” or time 0. In the Zespri example, fruit is harvested at this harvest time each year. The selling season in the destination market begins at time t_0 . During this season, the selling price of the product declines over time. We divide the season into $N + 1$ time periods, as illustrated in Figure 1. Let $p(t)$ denote the market price in time period t , where $t = 1, 2, \dots, N + 1$. Period $N + 1$ represents the salvage period with price $p(N + 1)$. The price schedule $p(t)$ is exogenously given and (weakly) decreasing in t . Without loss of generality, we assume that each period lasts one time unit (e.g., one day) and N is large enough to accommodate delayed shipments. That is, period t refers to time segment $[t_0 + t - 1, t_0 + t)$ for $t = 1, 2, \dots, N + 1$.

In the destination market, demand in period t ($t = 1, 2, \dots, N$) is random, denoted by $d(t)$. All leftover product can be salvaged in period $N + 1$. The $d(t)$ ’s need not be independent and their distributions may depend on the prices $p(1), p(2), \dots, p(N)$ as well as the quality of the product in period t . In this paper, however, we do not postulate any specific relation between demand and price but assume that they are exogenously given. For convenience, let $D(t) = \sum_{k=1}^t d(k)$, representing the total demand up to period t . $D(t)$ has a cumulative distribution function (cdf) $F_t(x)$. Let $\bar{F}_t(x) = 1 - F_t(x)$ and $f_t(x)$ denote the probability density function of $D(t)$. Unsatisfied demand in each period is backlogged, and can be filled in future periods but at a lower price depending on the actual arrival time. An orange exporter revealed in an interview that its customers normally accept late shipments but the selling price is adjusted according to the actual arrival time as a compensation for arriving late.

There are M potential vessels offering shipping services from the origin to the destination. For $i = 1, 2, \dots, M$, vessel i has a given departure time s_i , and will charge the shipper a freight rate r_i for each unit shipment.¹ For notational brevity, we discretize vessels’ arrival times as follows. Vessel i will arrive at the beginning of period T_i (or equivalently, at time $t_0 + T_i - 1$), where T_i is a discrete random variable with support $\{1, 2, \dots, N + 1\}$. The transit time of vessel i is hence given




Figure 1 Timeline and market markdown

by $t_0 + T_i - 1 - s_i$. The cdf of T_i is denoted by G_i , and the T_i 's need not be independent. In reality, the arrival times of different vessels are often correlated. For example, vessels arriving at the same port terminal will experience the same port congestion (if any). Since this paper focuses on ocean shipping which usually takes a much longer time than inland transportation, we assume that the time taken for inland distribution is constant and hence can be normalized to zero. Consequently, shipments are available for sales immediately after arrival.

Without loss of generality, we assume that $E[T_1] \leq E[T_2] \leq \dots \leq E[T_M]$, i.e., vessels are indexed in accordance with their expected arrival sequence. However, the freight rate r_i is not necessarily decreasing in i .² Among the M vessels in the model, some may belong to the same shipping service but depart at different times. Ocean carriers in reality usually provide weekly service schedules. For example, if the decision maker wants to take into consideration two possible departure times of a weekly shipping service, given that one period equals one day, we can treat these two options as vessels i and j with $s_j = s_i + 7$, $E[T_j] = E[T_i] + 7$ and $r_i = r_j$. Vessel j is then said to have a late departure.

The fruit is procured from growers at harvest time, and hence the shipper starts paying the inventory cost for each unit of product from time 0. Let h^O denote the inventory holding cost per period at the origin, including not only the capital cost but also the cost of maintaining a proper temperature in the warehouse. Let h_i^B be the holding cost rate on board a vessel, which may be vessel-dependent, as the difficulty in controlling the storage temperature for fruits often differs among chartered vessels and ocean liners. In the general model, we allow the shipper to hold leftover inventory at the destination to fulfill the demand in later periods. We denote by h^D the cost of holding one unit of product at the destination, which is incurred at the end of each period

Let q_i denote the quantity that the shipper decides to ship via vessel i . Then, the shipper's problem is to determine the allocation $\mathbf{q} = (q_1, q_2, \dots, q_M)$ in order to maximize the expected profit. As we are considering a single type of product and they are all purchased at harvest time, the unit procurement cost can, without loss of generality, be normalized to zero. Because ocean transport takes time, shipments are usually initiated long (e.g., 5-6 weeks) before the sale begins. As a result, the shipper is uncertain about the $d(t)$'s at the time of decision making, even though some vessels have late departure times. That is, our framework is an offline decision model in which we implicitly assume that the benefit of waiting for updated demand information with late departures is negligible. This is a reasonable assumption/approximation when the transit time is long compared to the lifetime of the product, as in the case of ocean shipping considered in this paper. In other cases (e.g., air freight transport), however, an online decision process would be a better fit.

Given $\mathbf{q}$, the shipper's income during period 1 includes the revenue at price $p(1)$ minus the cost of holding leftover inventory, which can be written as

$$\begin{aligned} & p(1) \min\{d(1), \sum_{i=1}^M q_i I(T_i = 1)\} - h^D [\sum_{i=1}^M q_i I(T_i = 1) - d(1)]^+ \\ &= p(1)d(1) - p(1)[d(1) - \sum_{i=1}^M q_i I(T_i = 1)]^+ - h^D [\sum_{i=1}^M q_i I(T_i = 1) - d(1)]^+, \end{aligned}$$

where $(x)^+ = \max\{x, 0\}$ and $I(\cdot)$ is an indicator function. The term $\sum_{i=1}^M q_i I(T_i = 1)$ represents the quantity that has arrived in period 1. The above equality follows by invoking the relation $\min\{x, y\} = x - (x - y)^+$.

In period 2, the accumulated demand is given by $[d(1) - \sum_{i=1}^M q_i I(T_i = 1)]^+ + d(2)$, and the total available inventory equals $[\sum_{i=1}^M q_i I(T_i = 1) - d(1)]^+ + \sum_{i=1}^M q_i I(T_i = 2)$. The shipper's income in the second period can then be written as

$$\begin{aligned} & p(2) \min\{[d(1) - \sum_{i=1}^M q_i I(T_i = 1)]^+ + d(2), [\sum_{i=1}^M q_i I(T_i = 1) - d(1)]^+ + \sum_{i=1}^M q_i I(T_i = 2)\} \\ & - h^D [\sum_{i=1}^M q_i I(T_i \leq 2) - D(2)]^+ \\ &= p(2)d(2) + p(2)[d(1) - \sum_{i=1}^M q_i I(T_i = 1)]^+ - p(2)[D(2) - \sum_{i=1}^M q_i I(T_i \leq 2)]^+ - h^D [\sum_{i=1}^M q_i I(T_i \leq 2) - D(2)]^+ \end{aligned}$$

Accordingly, the income in period t can be expressed as

$$p(t)d(t) + p(t)[D(t-1) - \sum_{i=1}^M q_i I(T_i \leq t-1)]^+ - p(t)[D(t) - \sum_{i=1}^M q_i I(T_i \leq t)]^+ - h^D [\sum_{i=1}^M q_i I(T_i \leq t) - D(t)]^+.$$

In period $N + 1$, both the overage quantity and the quantity arriving after period N will be salvaged. The shipper thus obtains

$$\begin{aligned} & p(N+1) \left\{ \left[\sum_{i=1}^M q_i I(T_i \leq N) - D(N) \right]^+ + \sum_{i=1}^M q_i I(T_i = N+1) \right\} \\ &= p(N+1) \left[D(N) - \sum_{i=1}^M q_i I(T_i \leq N) \right]^+ - p(N+1) D(N) + p(N+1) \sum_{i=1}^M q_i, \end{aligned}$$

where we have used the relation $(x - y)^+ = (y - x)^+ - (y - x)$.

Summing these $N + 1$ terms yields the total profit collected in the destination market, which equals $\sum_{t=1}^N (p(t) - p(N+1))d(t) - \sum_{t=1}^N b(t)[D(t) - \sum_{i=1}^M q_i I(T_i \leq t)]^+ - \sum_{t=1}^N h^D [\sum_{i=1}^M q_i I(T_i \leq t) - D(t)]^+ + p(N+1) \sum_{i=1}^M q_i$, where we define $b(t) = p(t) - p(t+1)$ for $t = 1, 2, \dots, N$. The first term is irrelevant to the optimization and can be omitted. In addition to the shipping cost and the inventory cost incurred before arrivals, the shipper's objective is to minimize the following expected cost function.

$$\text{Problem (1): } \min_{\mathbf{q} \geq 0} H(\mathbf{q}) = \sum_{t=1}^N \left\{ b(t) E[D(t) - \sum_{i=1}^M q_i I(T_i \leq t)]^+ + h^D E[\sum_{i=1}^M q_i I(T_i \leq t) - D(t)]^+ \right\} + \sum_{i=1}^M c_i q_i, \quad (1)$$

where we define

$$c_i = r_i + s_i h^O + (t_0 + E[T_i] - 1 - s_i) h^B - p(N+1).$$

When one unit of shipment is assigned to vessel i , the shipper will incur a shipping cost r_i , an inventory holding cost $s_i h^O$ at the origin, as well as the expected pipeline inventory cost $h^B(t_0 + E[T_i] - 1 - s_i)$. We will henceforth refer to c_i as the *full variable cost* of vessel i . As in the classical newsvendor problem, we assume that $c_i > 0$ for all i 's to rule out the trivial case in which the shipper would ship as much as possible.

Finally, the framework proposed here in general applies to two different scenarios, depending on who, the shipper or local distributors, is managing the downstream inventory.

Single-Demand Model. While holding inventory at the destination is allowed in the general framework, in some cases the shipper may not have any warehouse at the destination. Instead, local distributors would purchase the product at the beginning of the season and then manage the inventory themselves. To model this scenario, we assume that all customers will place their orders in period 1; as before, unsatisfied demand can be filled in future periods by later shipments, but the selling price is discounted according to the actual arrival time. Then, the total demand over the entire season can be modeled as a single random variable, denoted by D , which is realized in period 1. That is, $d(1) = D$ and $d(t) = 0$ for $t \geq 2$ (or equivalently, $D(t) = D$ for all $t = 1, 2, \dots, N$). Since

there is no demand late in the season, when D is fully satisfied, excess shipments will simply be salvaged after arrival. This is equivalent to setting $h^D = 0$. This scenario corresponds to a special case of our general framework in which $D(t) = D$ for all $t = 1, 2, \dots, N$ and $h^D = 0$. We will refer to this case as the single-demand model. In related discussions, we will denote the cdf of D by $F(x)$ and define $\bar{F}(x) = 1 - F(x)$.

Multi-Demand Model. If the shipper has a warehouse at the destination, customers may purchase the product at any time throughout the season as modeled in our general framework. We will refer to this case as the multi-demand model. Note that it is often cheaper to hold inventory at the origin than in the overseas market (i.e., $h^B < h^D$), as the warehouse is often well-established at the origin where the shipper can enjoy economies of scales. In the multi-demand model, consequently, it becomes especially important to take into account different departure times of the same shipping service, as late departures may be preferred in order to properly allocate the inventory costs incurred at the origin and the destination.

4. The Single-Demand Model with $M = 2$

We start by analyzing the single-demand model with $M = 2$. From this special case, we will demonstrate that Problem (1) combines two classes of newsvendor models in the literature.

With $D(t) = D$ for all t and $h^D = 0$, and Problem (1) can be written as

$$\min_{\mathbf{q} \geq 0} H(\mathbf{q}) = \sum_{t=1}^N \left\{ b(t) E \left[D - \sum_{i=1}^2 q_i I(T_i \leq t) \right]^+ \right\} + \sum_{i=1}^2 c_i q_i. \quad (2)$$

If both T_1 and T_2 are deterministic with $T_1 < T_2$, (2) can be simplified as

$$\min_{\mathbf{q} \geq 0} H(\mathbf{q}) = (p(1) - p(T_1)) E[D] + (p(T_1) - p(T_2)) E[D - q_1]^+ + (p(T_2) - p(N+1)) E[D - (q_1 + q_2)]^+ + \sum_{i=1}^2 c_i q_i. \quad (3)$$

Let us then express the expected markdown loss as the objective value of a linear program. $H(\mathbf{q})$ can then be rewritten as

$$H(\mathbf{q}) = \sum_{i=1}^2 c_i q_i + E \left\{ \begin{array}{l} \min \sum_{i=1}^2 w_i x_i + s x_s \\ s.t. \quad 0 \leq x_i \leq q_i, \forall i \\ \sum_{i=1}^2 x_i + x_s = D \\ 0 \leq x_s. \end{array} \right\}, \quad (4)$$

where $w_1 = p(1) - p(T_1)$, $w_2 = p(1) - p(T_2)$ and $s = p(1) - p(N+1)$.

Note that by assuming that both T_i are *deterministic*, we have reduced Problem (1) to a formulation equivalent to that studied in the option contract procurement literature (see, for example, Section 5 of Martínez-de Albéniz and Simchi-Levi (2005) and Section 4 of Martínez-de Albéniz and Simchi-Levi (2009)). In that stream of literature, w_i is referred to as the execution cost of

option contract i and s is regarded as the opportunity cost due to lost sales or a constant spot market price. q_i is the reservation capacity for option i , which must be determined before demand is realized. One can think of each vessel as a supply option contract: c_i is the reservation cost for option contract i ; the resulting selling price $p(T_i)$ is the marginal profit that can potentially be obtained by executing option contract i . Our shipping problem *with deterministic arrival times* can then be interpreted as a supply option contract procurement problem with $w_1 \leq w_2 \leq s$. The inner linear program of (4) can then be solved in a greedy fashion: Given realized demand D , we should clearly execute option 1 first, and then execute option 2, until D is fully filled. This leads exactly to the expected markdown loss expressed in (3).

Due to uncertain arrival times, a critical characteristic of ocean transport, our shipping problem departs from the existing option contract procurement models in that unreliable service schedules result in *random* execution costs w_i . To our knowledge, this portfolio selection problem with random reservation profits/costs has not been addressed in the literature. One relevant paper is Fu et al. (2010) in which the authors study how to deal with a random spot price s while assuming the w_i 's are deterministic.

Now assume that in (2), both T_1 and T_2 are generally random but $N = 1$. In this case, shipments will be sold either at the full price $p(1)$ or at the salvage value $p(2)$, depending on whether the arrival is delayed to the salvage period. Our model then coincides with the supply diversification (e.g., Anupindi and Akella 1993, Federgruen and Yang 2008) problem with 0-1 random yield. With probability $G_i(t_1)$, shipments via vessel i can sell at the full price; with probability $1 - G_i(t_1)$, they will miss the full-price sales window and must be salvaged.

In our shipping context, the shipper can still receive a partial profit margin if the delay is not too extreme, permitting a trade-off between speed and markdown loss. Hence, our model with $N \leq 2$ departs from the supply diversification literature by further specifying the length of delivery time and its impact on the shipper's cost, thus incorporating the speed-cost trade-off.³

To summarize, we have shown that our shipping problem can be reduced to two classes of newsvendor models from different perspectives. These two classes have been studied separately in the literature, and it appears that there are no models/applications in which they are combined.

4.1. Deterministic Arrival Sequence

Let us first consider a special case in which $\Pr\{T_1 \leq T_2\} = 1$, i.e., the vessels arrive according to a deterministic sequence. Theorem 1 characterizes the structure of the optimal solution under the assumption of a deterministic arrival sequence, suggesting that despite uncertainties in arrival times, one can simply replace random prices $p(T_1)$ and $p(T_2)$ with their expectations if $\Pr\{T_1 \leq T_2\} = 1$ and the structure of our problem therefore mirrors that of the supply option contract problem studied in the literature. All proofs in this paper are presented in Appendix EC.1.

THEOREM 1. Suppose $\Pr\{T_1 \leq T_2\} = 1$. The optimal solution $\mathbf{q}^*$ to (2) is given by one of the following cases:

- (i) If $c_i \geq E[p(T_i)] - p(N+1)$ for both $i = 1, 2$, $\mathbf{q}^* = (0, 0)$, i.e., it is optimal to ship nothing.
- (ii) If $\frac{c_1}{E[p(T_1)] - p(N+1)} \leq \frac{c_2}{E[p(T_2)] - p(N+1)}$ and $c_1 < E[p(T_1)] - p(N+1)$, then $\mathbf{q}^* = (q_1^s, 0)$ where $q_1^s = \bar{F}^{-1}\left(\frac{c_1}{E[p(T_1)] - p(N+1)}\right)$.
- (iii) If $E[p(T_1)] - c_1 \leq E[p(T_2)] - c_2$ and $c_2 < E[p(T_2)] - p(N+1)$, then $\mathbf{q}^* = (0, q_2^s)$ where $q_2^s = \bar{F}^{-1}\left(\frac{c_2}{E[p(T_2)] - p(N+1)}\right)$.
- (iv) If

$$E[p(T_1)] - c_1 > E[p(T_2)] - c_2, \text{ and} \quad (5)$$

$$\frac{c_1}{E[p(T_1)] - p(N+1)} > \frac{c_2}{E[p(T_2)] - p(N+1)}, \quad (6)$$

then the diversification solution (q_1^d, q_2^d) is optimal, where $q_1^d = \bar{F}^{-1}\left(\frac{c_1 - c_2}{E[p(T_1)] - E[p(T_2)]}\right)$ and $q_2^d = \bar{F}^{-1}\left(\frac{c_2}{E[p(T_2)] - p(N+1)}\right) - q_1^d$.




Figure 2 Optimal shipping strategies in the single-demand model with $M = 2$ and $\Pr\{T_1 \leq T_2\} = 1$

In Theorem 1, all possible combinations of nonnegative variable costs c_1 and c_2 are segmented into four regions according to the conditions in cases (i)-(iv). The segmentation is illustrated in Figure 2. In region (I), full variable costs are too high for both vessels such that the shipper would ship nothing. In regions (II) and (III), choosing only one vessel is optimal since one vessel is significantly cheaper than the other, where q_i^s denotes the shipping quantity when the shipper exclusively uses




Figure 3 Graphic explanation of Theorem 1

vessel i . In region (IV) where parameters satisfy inequalities (5) and (6), the shipper should use both vessels, where q_i^d denotes the shipping quantity assigned to vessel i .

As in the supply option contract problem where the buyer may need to maintain a portfolio of option contracts, the shipper uses a combination of different shipping modes to manage demand uncertainty. Inequality (5) clearly implies that vessel 1 generates a higher marginal profit than vessel 2. To interpret inequality (6), note that when the shipper exclusively uses vessel i , the solution is determined by $\bar{F}(q_i^s) = \frac{c_i}{E[p(T_i)] - p(N+1)}$, as in the classical newsvendor problem. That is, inequality (6) implies that vessel 2, if used exclusively, will lead to a higher (type 1) service level than vessel 1. A graphic explanation is given in Figure 3. Diversification is optimal if the fast vessel provides a higher marginal profit and the slow one yields a higher service level. The early arrival enables shipments to be sold at a premium, whereas the shipments arriving late contribute to fulfilling more demand in a probabilistic sense.

COROLLARY 1. *The optimal shipping quantities in Theorem 1 have the following properties:*

- (i) $q_1^d + q_2^d = q_2^s$, i.e., the total shipping quantity with diversification depends only on the parameters of the slow vessel.
- (ii) If $\hat{D}$ is a mean-preserving spread of D , then the corresponding solutions satisfy $q_2^d \leq \hat{q}_2^d$, $|q_1^d - \mu_D| \leq |\hat{q}_1^d - \mu_D|$ and $|q_1^d + q_2^d - \mu_D| \leq |\hat{q}_1^d + \hat{q}_2^d - \mu_D|$. In particular, as $\text{Var}(D) \rightarrow 0$, $q_1^d \rightarrow \mu_D$ and $q_2^d \rightarrow 0$.
- (iii) Suppose $h_1^B = h_2^B = h^B$. If service i ($i = 1$ or 2) has a shorter transit time, then q_i^d is increasing in h^B , whereas q_{3-i}^d is decreasing in h^B .

Part (i) of Corollary 1 states that when it is optimal to diversify shipments, the shipper should order up to the same level as it does when only using the slow vessel, but assign a portion of

shipments to the fast vessel in order to enjoy a higher profit margin. Part (ii) of Corollary 1 reveals the effect of demand variability on shipping strategies. The quantity shipped via the slow vessel is increasing in demand variability, whereas the quantity shipped via the fast vessel and the total quantity will approach the mean of demand as demand variability decreases. In the extreme case where demand is known accurately, it boils down to a simple principle: Select the vessel that would give the highest profit for every unit of shipments. However, when demand uncertainty exists, the cheaper but slower vessel may be used as a buffer against demand variability. Part (iii) of Corollary 1 reflects the intuitive trade-off between the pipeline inventory cost and the shipping cost. When holding pipeline inventory is more expensive, a larger portion of shipments should be assigned to the fast vessel.

4.2. Uncertain Arrival Sequence

The preceding discussion has relied on the assumption that $T_1 \leq T_2$ with probability 1, which excludes another portfolio effect with respect to the shipping schedule reliability. Theorem 2 presents the structure of optimal shipping strategies when this assumption is relaxed.

THEOREM 2. *For $i = 1, 2$, the optimal solution to (2) is characterized as follows:*

- (i) *If $c_i \geq \sum_{t=1}^N b(t)G_i(t)$ for both $i = 1, 2$, then $\mathbf{q}^* = (0, 0)$, i.e., it is optimal to ship nothing;*
- (ii) *For $i = 1, 2$, if $c_{3-i} - \frac{\sum_{t=1}^N b(t)Pr\{T_i \leq t, T_{3-i} \leq t\}}{\sum_{t=1}^N b(t)G_i(t)}c_i - \sum_{t=1}^N b(t)Pr\{T_i > t, T_{3-i} \leq t\} \geq 0$ and $c_i < \sum_{t=1}^N b(t)G_i(t)$, then $q_i^* = q_i^s$, $q_{3-i}^* = 0$, where q_i^s is determined by*

$$\bar{F}(q_i^s) = \frac{c_i}{\sum_{t=1}^N b(t)G_i(t)}; \quad (7)$$

- (iii) *Otherwise, $\mathbf{q}^* = (q_1^d, q_2^d)$, which is obtained by solving the following equations:*

$$\begin{aligned} c_i = & \bar{F}(q_1^d + q_2^d) \sum_{t=1}^N b(t)Pr\{T_i \leq t, T_{3-i} \leq t\} \\ & + \bar{F}(q_i^d) \sum_{t=1}^N b(t)Pr\{T_i \leq t, T_{3-i} > t\}, \text{ for } i = 1, 2. \end{aligned} \quad (8)$$

The structure of the optimal solution is illustrated in Figure 4, which mirrors some existing results in the supply risk diversification literature.⁴ Compared with Figure 2, the region for diversification is enlarged from a triangle to a quadrangle, as the shipper in this case also needs to diversify its shipments in order to mitigate the delay risk.

One may also view each vessel as an option contract. When vessels arrive in a random order, $p(T_1)$ may be either higher or lower than $p(T_2)$. Just like investing in a portfolio of assets with random returns, when deciding on the shipping/reservation quantities, the shipper can use multiple vessels/option contracts to control the volatility of their execution profits. Therefore, we conclude




Figure 4 Optimal shipping strategies in the single-demand model with $M = 2$

that in the shipping context, the benefits of using multiple vessels are twofold: Managing demand uncertainty via diversifying the timing of sales, and mitigating the delay risk via splitting shipments among vessels.

REMARK 1. We are now able to position our shipping problem (1) in the extensive literature on the newsvendor procurement. Problem (1) provides a unified view of the procurement problem with supply risk diversification and supply option contracts in the sense that it generalizes the supply diversification problem with 0-1 yield from the perspective of shipping reliability, and the supply option contract selection from the perspective of managing demand uncertainty.

4.3. Diversification with Similar or Different Services?

In this subsection, we provide a numerical example to examine the interplay between the two portfolio effects. For simplicity, let $t_0 = 0$ (i.e., demand at the destination begins accruing at time 0), and $N = 14$, $p(t) = 61 - \frac{30}{7}(t-1)$ for $t = 1, 2, \dots, N+1$, $h^O = h_i^B = 0.1$ for all i . D follows a gamma distribution with a mean $\mu_D = 100$, whose coefficient of variation is denoted by cv_D . Consider four vessels all departing at time 0. The distributions of T_i 's are constructed from normally distributed random variables with means $\mu_{T1} = \mu_{T2} = 6$, $\mu_{T3} = \mu_{T4} = 8.5$ and the same coefficient of variation cv_T . The coefficient of correlation, denoted by ρ , is set to be identical among all underlying arrival times. The first two fast vessels are associated with freight rates $r_1 = r_2 = 20$, whereas the other two with $r_3 = r_4 = 10$.

The question is whether to choose similar arrival schedules or different ones. A portfolio of vessels with different arrival schedules is preferred in order to tackle the demand-side uncertainty. However, the diversification effect on mitigating delay risk is somewhat weakened, because the arrival sequence is less likely to alter if the arrival schedules differ significantly. On the other hand,

leveraging two similar arrival schedules strengthens the diversification effect on schedule reliability, but would lessen the effect on demand uncertainty.




Figure 5 Comparison of two-vessel portfolios

We compare the performances of two-vessel portfolios. Three possible strategies are possible: FF (selecting both fast services), FS (selecting heterogeneous services), and SS (selecting both slow services). With different pairs of cv_D and cv_T , the best strategies are shown in Figure 5a. When cv_T is small, using heterogeneous services is always the best strategy. FF outperforms FS when cv_D is low but cv_T is relatively high, since in this case spreading delay risk is of a higher priority. Nevertheless, it should also be kept in mind that FS does spread some delay risk if the vessels' arrival sequence can be altered. That is why the region for FF starts shrinking when cv_T exceeds 0.2. When both demand and arrival times are highly volatile, SS becomes the best strategy. Figure 5b presents the results with different cv_D and ρ . For a small ρ , it is more beneficial to use two identical vessels to mitigate the delay risk when cv_D is either small or extremely large. However, FS becomes dominant gradually as ρ increases. Even though the correlation between arrival times lessens the effect that diversification has on spreading the delay risk, the strategy FS remains efficacious in dealing with demand uncertainty.

5. The General Case: $M > 2$

Despite the convexity of Problem (1), it remains difficult to solve because evaluating the expectation over $D(t)$ and the T_i 's involves many integrals, which are computationally demanding. In this section, we address Problem (1) step by step. In Section 5.1, we analyze the special case where vessels' arrival sequence is deterministic. The appealing structure in this case allows us to derive the optimal solution in a closed form (in the single-demand model) or numerically solve it in polynomial

time (in the multi-demand model). In Section 5.2, we propose an iterative approximation procedure, leveraging the results in Section 5.1, to address the general problem with an uncertain arrival sequence.

5.1. Deterministic Arrival Sequence

In this subsection, we focus on a restrictive version of Problem (1) under the following assumption:

ASSUMPTION 1 (Deterministic Arrival Sequence). $\Pr\{T_i \leq T_{i+1}\} = 1$ for all $i = 1, 2, \dots, M-1$.

Assumption 1 is not entirely unrealistic. It is possible that all vessels are processed by the same port terminal operator at the destination. If the terminal operator must follow a predetermined service schedule according to the expected arrival times, Assumption 1 is then satisfied. Note that Assumption 1 has excluded the diversification effect on mitigating uncertain arrival schedules.

5.1.1. The Single-Demand Model In Theorem 1 and the accompanying discussions, we have identified a connection between our problem and the option contract procurement problem when $M = 2$. The following lemma extends this observation to the case when $M > 2$, indicating that uncertain arrivals affect the objective value only through the expected selling prices $\bar{p}_i = E[p(T_i)]$ as long as vessels' arrival sequence is fixed.

LEMMA 1. *In the single-demand model, if Assumption 1 is satisfied, $H(\mathbf{q})$ can be simplified to*

$$\begin{aligned} H(\mathbf{q}) = & (p(1) - \bar{p}_1)E[D] + \sum_{i=1}^{M-1} (\bar{p}_i - \bar{p}_{i+1})E[D - \sum_{j=1}^i q_j]^+ \\ & + (\bar{p}_M - p(N+1))E[D - \sum_{j=1}^M q_j]^+ + c_i q_i, \end{aligned} \quad (9)$$

where $\bar{p}_i = E[p(T_i)]$ for all i .

In view of (9), determining $\mathbf{q}$ can be regarded as equivalent to deciding on the reservation quantities of M supply option contracts where contract i , which corresponds to vessel i , has a reservation cost c_i and yields a random profit $p(T_i)$ from executing one unit of the reservation. Because under Assumption 1 we have $p(T_1) \geq p(T_2) \geq \dots \geq p(T_M)$ with probability one, the priority in executing these contracts is unambiguous. The shipper will first execute contract 1, and then contract 2, and so on until D is fully satisfied. This special structure enables us to explicitly characterize the optimal choice of $\mathbf{q}^*$ that minimizes $H(\mathbf{q})$.

We define

$$i_A = \arg \max_{1 \leq i \leq M} \{\bar{p}_i - c_i\}, \quad (10)$$

$$i_B = \arg \min_{1 \leq i \leq M} \left\{ \frac{c_i}{\bar{p}_i - p(N+1)} \right\}. \quad (11)$$

Hence, vessel i_A generates the highest marginal profit, and vessel i_B , if used exclusively, leads to the highest service level. In case of multiple maxima or minima, i_A and i_B are then selected as the lexicographically smallest.

We say vessel i is active if $q_i > 0$. Let J be the set of active vessels. The set J can be found by the following procedure.

ALGORITHM 1.

Step 0 Find i_A and i_B as defined in (10) and (11). Let $J := \{i_B\}$. If $i_A = i_B$, stop; otherwise, go to Step 1.

Step 1 $j_1 := \arg \min_{i_A \leq j < i_B} \left\{ \frac{c_j - c_{i_B}}{\bar{p}_j - \bar{p}_{i_B}} : \bar{p}_j - c_j \geq \bar{p}_{i_B} - c_{i_B} \right\}$. Let $J := J \cup \{j_1\}$. If $i_A = j_1$, stop; otherwise, go to step 2.

Step k (k=2,3, ...) $j_k := \arg \min_{i_A \leq j < j_{k-1}} \left\{ \frac{c_j - c_{j_{k-1}}}{\bar{p}_j - \bar{p}_{j_{k-1}}} : \bar{p}_j - c_j \geq \bar{p}_{j_{k-1}} - c_{j_{k-1}} \right\}$. Let $J := J \cup \{j_k\}$. If $i_A = j_k$, stop; otherwise, go to step $(k+1)$.




Figure 6 Illustration of the efficient frontier

THEOREM 3. Let $J = \{i_A, j_k, \dots, j_1, i_B\}$ be the set generated by Algorithm 1. The optimal solution $\mathbf{q}^*$ to the single-demand model under Assumption 1 is determined as follows. $q_i^* = 0$ whenever $i \notin J$. For $i \in J$,

(i) If $J = \{i_B\}$, then

$$q_{i_B}^* = \bar{F}^{-1} \left(\frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \right).$$

(ii) If $J = \{i_A, i_B\}$, then

$$\begin{aligned} q_{i_A}^* + q_{i_B}^* &= \bar{F}^{-1} \left(\frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \right), \\ q_{i_A}^* &= \bar{F}^{-1} \left(\frac{c_{i_A} - c_{i_B}}{\bar{p}_{i_A} - \bar{p}_{i_B}} \right). \end{aligned}$$

(iii) If $J = \{i_A, j_k, \dots, j_1, i_B\}$, then

$$\begin{aligned} q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^* + q_{i_B}^* &= \bar{F}^{-1} \left(\frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \right), \\ q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^* &= \bar{F}^{-1} \left(\frac{c_{j_1} - c_{i_B}}{\bar{p}_{j_1} - \bar{p}_{i_B}} \right), \\ q_{i_A}^* + \sum_{i=j_\tau}^{j_k} q_i^* &= \bar{F}^{-1} \left(\frac{c_{j_\tau} - c_{j_{\tau-1}}}{\bar{p}_{j_\tau} - \bar{p}_{j_{\tau-1}}} \right), \text{ for } \tau = 2, \dots, k, \\ q_{i_A}^* &= \bar{F}^{-1} \left(\frac{c_{i_A} - c_{j_k}}{\bar{p}_{i_A} - \bar{p}_{j_k}} \right). \end{aligned}$$

Algorithm 1 constructs an efficient frontier: Only the vessels on the frontier are assigned with a positive volume of shipments, as illustrated in Figure 6. Once the set $J = \{i_A, j_k, \dots, j_1, i_B\}$ is determined, $\mathbf{q}^*$ can be derived in a closed form, as given in Theorem 3. The efficient frontier and the solution structure mirrors the “lower convex envelope” property derived in the literature on supply option contract procurement (e.g., Martínez-de Albéniz and Simchi-Levi 2005, 2009). From the perspective of option contract procurement, Theorem 3 extends existing results by allowing for random execution profits/costs, indicating that the solution structure under deterministic execution profits can be sustained as long as random execution profits are always realized in a certain order. For our shipping application, this provides not only a building block for efficiently solving Problem (1), but also qualitative guidance for selecting shipping lines. The structure of the efficient frontier suggests that the shipper should try to obtain a high marginal profit via fast vessels, and meanwhile maintain as high a service level as possible by diversifying with slow vessels. It is also interesting to find that the overall order-up-to level is determined by $c_{i_B}/(\bar{p}_{i_B} - p(N+1))$, which depends only on vessel i_B . That is, the slowest active vessel helps maintain an overall service level, while the faster ones contribute to increasing profit margins.

COROLLARY 2. For $J = \{i_A, j_k, \dots, j_1, i_B\}$ generated by Algorithm 1, we have $c_{i_A} \geq c_{j_k} \geq c_{j_2} \geq c_{j_1} \geq c_{i_B}$.

Recall that different departures of every shipping line can be represented by different candidates in our model. However, Corollary 2 states that in the optimal portfolio, a vessel expected to arrive earlier would be associated with a higher variable cost, suggesting that we only need to consider the earliest departure of each shipping line as a potential option. Vessels with late departures must not be included in the optimal portfolio, since they arrive later but have larger c_i .

5.1.2. The Multi-Demand Model We proceed to study the multi-demand model in which the shipper may hold inventory at the destination at a holding cost rate h^D and allow downstream customers to order over time. As a result, contrary to Corollary 2 in the single-demand model, the shipper may choose some late departures to reduce the high inventory costs at the destination.

For ease of exposition, we define

$$\begin{aligned} L_0 &= E_T \left[\sum_{t=1}^{T_1-1} b(t)D(t) \right], \\ L_i(x) &= E_T \left[\sum_{t=T_i}^{T_{i+1}-1} b(t)(D(t)-x)^+ + h^D(x-D(t))^+ \right] \text{ for } i = 1, 2, \dots, M. \end{aligned} \quad (12)$$

One may interpret L_0 as the expected markdown loss before the arrival of vessel 1, and $L_i(x)$ as the expected markdown and inventory holding costs incurred between the arrivals of vessel i and vessel $i+1$ if the total shipping quantity up to vessel i equals x . Then, under Assumption 1, $H(\mathbf{q})$ can be simplified as follows:

$$\min_{\mathbf{q} \geq 0} H(\mathbf{q}) = L_0 + \sum_{i=1}^M L_i \left(\sum_{j=1}^i q_j \right) + \sum_{i=1}^M c_i q_i. \quad (13)$$

From the perspective of supply option procurement, the first term in (13) can be interpreted as a generalized cost function for executing reserved quantities q_i . One may think of T_i as the default execution date associated with option contract i . Contract i can be executed only when $t \geq T_i$. In every period t , as long as the cumulative demand $D(t)$ has not been completely filled, the shipper should execute the contracts available at period t and accrue a profit $p(t)$ for each unit; when $D(t)$ is fully satisfied, the shipper should defer the execution where the inventory holding cost h^D may be viewed as a penalty for late execution. Note that this portfolio selection problem, even with deterministic T_i , is a departure from the classical option contract problem in that the multi-demand setting gives rise to another trade-off, i.e., that between the inventory holding costs at the origin and the destination.

Let $J = \{j_1, j_2, \dots, j_{|J|}\}$ denote the set of active vessels in the optimal solution. We first state the first-order necessary optimality condition for $\mathbf{q}^*$ in the following lemma.

LEMMA 2. *Suppose that $J = \{j_1, j_2, \dots, j_{|J|}\}$ is the set of indices such that $q_i^* > 0$ for $i \in J$ and $q_i^* = 0$ for $i \notin J$. Under Assumption 1, the value of $\mathbf{q}^*$ must satisfy*

$$(\bar{p}_{j_k} - c_{j_k}) - (\bar{p}_{j_{k+1}} - c_{j_{k+1}}) = E \left[\sum_{t=T_{j_k}}^{T_{j_{k+1}}-1} (b(t) + h^D) F_t \left(\sum_{v=1}^k q_{j_v}^* \right) \right] \text{ for } k = 1, 2, \dots, |J| - 1, \quad (14)$$

$$\bar{p}_{j_{|J|}} - c_{j_{|J|}} - p(N+1) = E \left[\sum_{t=T_{j_{|J|}}}^N (b(t) + h^D) F_t \left(\sum_{v=1}^{|J|} q_{j_v}^* \right) \right]. \quad (15)$$

Once J is determined, $\mathbf{q}^*$ can be computed according to Lemma 2. However, there are many possible subsets of $\{1, 2, \dots, M\}$ constituting a candidate solution. We need to identify the one that generates the minimum cost. An attractive property of the shipping quantities is as follows: The

cumulative quantity up to vessel j_k is determined only by the parameters of vessels j_k and j_{k+1} ; the total quantity is determined only by the parameters of the last (i.e., slowest) vessel in the portfolio. If we work with the cumulative quantities $y_i = \sum_{j=1}^i q_j$, then y_i is nondecreasing in i . Provided that $J = \{j_1, j_2, \dots, j_{|J|}\}$, the optimal cumulative quantities are given by $y_{j_k}^* = \sum_{v=1}^k q_{j_v}^*$ for all $k = 1, 2, \dots, |J|$. Note that $y_{j_k}^*$ is *strictly* increasing in k . Accordingly, we rewrite the objective function of (13) as

$$\begin{aligned} H(\mathbf{y}) &= L_0 + \sum_{i=1}^M L_i(y_i) + \sum_{i=2}^M c_i(y_i - y_{i-1}) + c_1 y_1 \\ &= L_0 + \sum_{i=1}^{M-1} \{L_i(y_i) + (c_i - c_{i+1})y_i\} + L_M(y_M) + c_M y_M, \end{aligned}$$

where we have regrouped the terms such that the total cost can be written as a summation of terms with respect to y_i . Furthermore, given the optimal portfolio $J = \{j_1, j_2, \dots, j_{|J|}\}$, we have $y_i^* = 0$ for all $i < j_1$, $y_i^* = y_{j_k}^*$ for all $j_k \leq i \leq j_{k+1} - 1$ where $k = 1, 2, \dots, |J| - 1$, and $y_i^* = y_{j_{|J|}}^*$ for all $i \geq j_{|J|}$. The optimal value of the objective function is then reduced to

$$H(\mathbf{y}^*) = L_0 + w_{0,j_1} + \sum_{k=1}^{|J|-1} w_{j_k, j_{k+1}}(y_{j_k}^*) + w_{j_{|J|}, M+1}(y_{j_{|J|}}^*), \quad (16)$$

where we define

$$w_{0,i} = \sum_{k=1}^{i-1} L_k(0) \quad (17)$$

$$w_{i,j}(y) = \sum_{k=i}^{j-1} L_k(y) + (c_i - c_j)y \quad (18)$$

$$w_{i,M+1}(y) = \sum_{k=i}^M L_k(y) + c_i y \quad (19)$$

In (16), the constant term L_0 is irrelevant to the optimization and the other terms are expressed by $|J| + 1$ parts, enabling us to evaluate the cost incurred by activating one more vessel.

Consider a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$. Let the set of vertices $\mathcal{V} = \{O, 1, 2, \dots, M, E\}$: Vertices $i = 1, 2, \dots, M$ denote M candidate vessels and two artificial vertices O and E are added as the origin and end vertices. For edges, construct edge (i, j) by joining vertices i and j for any i, j such that $1 \leq i < j \leq M$. Also, connect vertex O with all other vertices as edges (O, i) for all $1 \leq i \leq M$ and (O, E) , and connect each $i = 1, \dots, M$ with the end vertex E as the edge (i, E) . Associated with each edge, there are two parameters y_{ij} and d_{ij} , keeping track of the aggregate quantity and costs respectively. The parameters are defined below and Figure 7 gives an illustration of graph $\mathcal{G}$.

- For edge (i, j) where $1 \leq i < j \leq M$, let y_{ij} be the solution to

$$(\bar{p}_i - c_i) - (\bar{p}_j - c_j) = E\left[\sum_{t=T_i}^{T_j-1} (b(t) + h^D) F_k(y_{ij})\right], \quad (20)$$

and set $y_{ij} = 0$ if no solution exists. Let $d_{ij} = w_{i,j}(y_{ij})$.

- For edge (O, i) where $1 \leq i \leq M$, let $y_{O,i} = 0$ and $d_{O,i} = w_{0,i}$.
- For edge (i, E) where $1 \leq i \leq M$, let $y_{i,E}$ be the solution to

$$\bar{p}_i - c_i - p(N+1) = E\left[\sum_{t=T_i}^N (b(t) + h^D) F_k(y_{i,E})\right], \quad (21)$$

and set $y_{i,E} = 0$ if no solution exists. Let $d_{i,E} = w_{i,M+1}(y_{i,E})$.

- For edge (O, E) , let $y_{O,E} = 0$ and $d_{O,E} = \sum_{i=1}^M L_i(0)$.




Figure 7 Illustration of graph $\mathcal{G}$ ($1 \leq i < j \leq M$)

A path is called a *monotone path* in $\mathcal{G}$ if it starts from vertex O and ends at vertex E with y_{ij} strictly increasing along the path. From Lemma 2, y_{ij} corresponds to the cumulative quantity if vessels i and j are consecutive in J , and $y_{i,E}$ is the total quantity if vessel i is the last arriving vessel in J . Thus, any feasible portfolio can be represented by a monotone path in $\mathcal{G}$. On the other hand, a path from O to E that passes $|J|$ vessels contains $|J| + 1$ edges. Note that the distances d_{ij} of these edges correspond to the $|J| + 1$ parts in (16). That is, the total distance of a monotone path equals the expected cost (minus a constant L_0) incurred by the corresponding solution. Therefore, we have established the following theorem.

THEOREM 4. *Under Assumption 1, determining the set J is equivalent to finding the shortest monotone path in $\mathcal{G}$.*

This variation of shortest-path problems can be solved by a modified shortest-path algorithm which is detailed in Appendix EC.2. Given the parameters (y_{ij}, d_{ij}) , the complexity of finding the monotone shortest path is $O(M^3)$. We therefore conclude that under Assumption 1, $\mathbf{q}^*$ can also be determined efficiently under multiple demands.

REMARK 2. The monotone shortest path algorithm is reminiscent of Fu et al. (2010) in which the authors develop an algorithm to solve the option contract procurement problem described in (4) when the spot market price s and demand D are interdependent random variables and execution costs w_i are deterministic. To some extent, our problem has a similar flavor: The cumulative demand $D(t)$ is correlated with selling price $p(t)$ through time t . However, our multi-demand model is different from that studied in Fu et al. (2010) in that multiple demands occur over time in our model which leads to another trade-off between inventory holding costs at the origin and the destination. Interestingly, despite the differences we have identified a similar property to that in Fu et al. (2010), i.e., the cumulative order quantities are determined only by the two consecutively active candidates in the portfolio and the total quantity depends only on the last active candidate, which enables the portfolio selection to be converted into a variation of the shortest-path problem.

5.2. Uncertain Arrival Sequence

Having discussed the solution approaches under Assumption 1, we are ready to propose an iterative approximation procedure to solve the general problem based on the findings in Section 5.1.

We first approximate the cost function by replacing random T_i 's with a fixed *nominal* arrival schedule: $\hat{T}_1 < \hat{T}_2 < \dots < \hat{T}_M < \hat{T}_{M+1} = N + 1$. Note that we require strict inequalities here for the sake of convergence (as will be discussed later). One can simply set $\hat{T}_i = E[T_i]$ if vessels' expected arrival times are distinct from each other.

As in (12), we define a set of functions to represent the costs incurred between the arrivals of every two consecutive vessels except that the uncertain T_i 's are replaced by $\hat{T}_i$'s:

$$\begin{aligned}\hat{L}_0 &= \sum_{t=1}^{\hat{T}_1} b(t)D(t), \\ \hat{L}_i(x) &= \sum_{t=\hat{T}_i}^{\hat{T}_{i+1}-1} (b(t)(D(t) - x)^+ + h^D(x - D(t))^+).\end{aligned}$$

The approximate objective function, denoted by $\hat{H}^0(\mathbf{q})$, is written as

$$\hat{H}^0(\mathbf{q}) = \hat{L}(\mathbf{q}) + \sum_{i=1}^M c_i q_i. \quad (22)$$

where $\hat{L}(\mathbf{q}) = \hat{L}_0 + \sum_{i=1}^M \hat{L}_i(\sum_{j=1}^i q_j)$.

Clearly, to minimize $\hat{H}^0(\mathbf{q})$ over $\mathbf{q} \geq \mathbf{0}$, we can simply apply the results in Section 5.1. Let $\mathbf{q}^0$ be the minimizer of this simple approximate problem. To improve the approximation, we can utilize the linear term $(\nabla H(\mathbf{q}^0) - \nabla \hat{H}(\mathbf{q}^0))^T \mathbf{q}$, which captures the difference between the gradients of the actual cost function and its approximation at $\mathbf{q}^0$. Inspired by the stochastic gradient-based

approximation, we can add this linear term multiplied by a proper stepsize/smoothing factor δ_0 to the initial approximation. By doing so, we are in effect adjusting the original variable cost c_i to $c_i + \delta_0(\nabla_{q_i} H(\mathbf{q}^0) - \nabla_{q_i} \hat{H}(\mathbf{q}^0))$ while preserving the appealing structure of $\hat{L}(\mathbf{q})$. With a given sequence of δ_k , we can continue this procedure to iteratively calibrate $\mathbf{c}$. Furthermore, note that conditioning on $\mathbf{T}$, the gradient of the true cost function $\nabla H(\mathbf{q}|\mathbf{T})$ has a closed-form expression and is thus easy to compute. Let $T_{(i)}$ denote the order statistics of $\mathbf{T}$ and (i) represent the index of the i -th arriving vessel. We have

$$\nabla_{(i)} H(\mathbf{q}|\mathbf{T}) = \sum_{k=i}^M \sum_{t=T_{(k)}}^{T_{(k+1)}-1} (b(t) + h^D) F_t\left(\sum_{j=1}^k q_{(j)}\right) - [p(T_{(i)}) - p(N+1)] + c_{(i)}.$$

With a set of independent random samples of arrival schedules, $\mathbf{T}(1), \mathbf{T}(2), \dots, \mathbf{T}(N_s)$, we can use $\frac{1}{N_s} \sum_{s=1}^{N_s} \nabla H(\mathbf{q}|\mathbf{T}(s))$ as an estimator of $\nabla H(\mathbf{q})$ in each iteration where N_s can be any positive integer. A minor point to note is that when using the stochastic gradient to calibrate $\mathbf{c}$, the updated variable cost may be so negative that the approximate problem becomes unbounded in some iteration (although this has never occurred in any of our numerical tests). To avoid this, we may introduce an artificial upper bound $\bar{\mathbf{Q}} = (\bar{Q}_1, \bar{Q}_2, \dots, \bar{Q}_M)$ for $\mathbf{q}$ where the $\bar{Q}_i$'s are sufficiently large numbers so that they will not affect the original optimal solution.⁵ The above discussion is formalized in Algorithm 2 below.

ALGORITHM 2.

Step 0 Let $\mathbf{c}^0 = \mathbf{c}$. Obtain $\mathbf{q}^0 = \arg \min_{\mathbf{0} \leq \mathbf{q} \leq \bar{\mathbf{Q}}} \hat{H}^0(\mathbf{q}) = \hat{L}(\mathbf{q}) + (\mathbf{c}^0)^T \mathbf{q}$.

Step k (k=1,2,...) Generate N_s random (and independent) samples of arrival schedules, $\mathbf{T}(1), \mathbf{T}(2), \dots, \mathbf{T}(N_s)$, and set

$$\mathbf{c}^k = \mathbf{c}^{k-1} + \delta_{k-1} \left(\frac{1}{N_s} \sum_{s=1}^{N_s} \nabla H(\mathbf{q}^{k-1}|\mathbf{T}(s)) - \nabla \hat{H}^{k-1}(\mathbf{q}^{k-1}) \right).$$

Obtain $\mathbf{q}^k = \arg \min_{\mathbf{0} \leq \mathbf{q} \leq \bar{\mathbf{Q}}} \hat{H}^k(\mathbf{q}) = \hat{L}(\mathbf{q}) + (\mathbf{c}^k)^T \mathbf{q}$. Go to Step $k+1$.

In general, Algorithm 2 falls into a class of stochastic approximation algorithms proposed by Cheung and Powell (2000). In that paper, the authors propose a generic approximation procedure intended for two-stage stochastic programs. However, they do not mention how to construct the initial approximation based on specific problems. Here, our algorithm has fully exploited the special structure given a deterministic arrival sequence in every iteration while calibrating the variable cost with the gradient information. In each iteration, $\hat{H}^k(\mathbf{q})$ can be efficiently optimized by applying the results in Section 5.1. The artificial upper bound $\bar{Q}_i$ has little impact here because it is attained only if $\lim_{q_i \rightarrow \infty} \nabla_i L(\mathbf{q}) + c_i^k = (N - \hat{T}_i + 1)h^D + c_i^k < 0$. This can be easily checked beforehand.

THEOREM 5. *If (i) for $k = 1, 2, \dots$, δ_k satisfies $0 < \delta_k < 1$, $\sum_{k=1}^{\infty} \delta_k = \infty$, $\sum_{k=1}^{\infty} \delta_k^2 < \infty$, and (ii) $\sum_{t=\hat{T}_i}^{\hat{T}_{i+1}-1} (b(t) + h^D)f_t(x) > 0$ for all $i = 1, 2, \dots, M$ and $x \in [0, \sum_{k=1}^i \bar{Q}_k]$, then the sequence of vectors $\mathbf{q}^k$ converges to an optimal solution $\mathbf{q}^*$ almost surely.*

Theorem 5 provides conditions under which Algorithm 2 will converge to an optimal solution. Condition (i) is common for the selection of step size. For instance, we may set $\delta_k = \frac{1}{k+1}$, which clearly satisfies condition (i). Condition (ii) guarantees that the approximate objective $\hat{H}^k$ is strongly convex, from which one can see that the nominal arrival schedules must be distinct from each other. Intuitively, if $\hat{T}_i = \hat{T}_{i+1}$, the deterministic arrival sequence suggests that selecting both vessels i and $i + 1$ is always unnecessary for approximate problems, but it can be optimal for the actual problem where the arrival sequence is uncertain. To avoid this degenerate case, if $E[T_i] = E[T_{i+1}]$ for some i , we may set $\hat{T}_i = E[T_i]$ and $\hat{T}_{i+1} = E[T_{i+1}] + 1$. Finally, the algorithm converges after an infinite number of iterations, which guarantees a near-optimal solution after sufficiently many iterations. Because the approximate problem in each iteration is solved in polynomial time with the methods proposed in Section 5.1, decision makers can easily control the total computational time to obtain a solution with reasonable quality.

6. Numerical Study

6.1. The Zespri Case

We set up numerical examples based on the Zespri case. The New Zealand company sells kiwifruit to many countries around the world, and Europe is one of its major destination markets. In New Zealand, kiwifruit ripens in early May every year. The maximum lifetime is 3-6 months.⁶ Europe is one of Zespri's major markets, and we will therefore consider Rotterdam as the destination port. Taking into account the time required for local distribution and downstream sales as well as the minimum time for the intercontinental transport, we assume that the selling season starts on the 20-th day after harvest ($t_0 = 20$) and lasts 70 days ($N = 70$). Available liner services from Port of Auckland to Port of Rotterdam with their transit times can be found at www.searates.com. The shortest transit time that a chartered vessel can provide may be calculated based on distance and sailing speed, which is around 24 days.⁷ Table 1 lists the expected transit times for seven representative shipping services. We further assume that transit times follow a multivariate normal distribution with means given in Table 1. For simplicity, the coefficients of variation are assumed identical for all i , and are denoted by cv_T . The correlation coefficients between pairs of the transit times are also identical, and are denoted by ρ . For simplicity, let every service offer weekly departures starting from time 0. Then, we can obtain the distributions of T_i 's in our model by appropriately discretizing the underlying arrival times. Based on the empirical evidence reported

Table 1 Representative shipping services from Auckland to Rotterdam

Shipping Service	Charter	SL 1	SL2	SL 3	SL 4	SL 5	SL 6
Expected Transit Time (days)	24	33	40	42	43	44	46

in Van de Weijer (2013), we set $\mathbf{r} = (28, 24, 18, 16, 15, 14, 12)$ such that the freight rate is decreasing in the transit time of each service.

In the fruit wholesale market, the selling price is affected by the product's condition such as the degree of maturity and decay that have occurred since the harvest (Agricultural Marketing Service 2012, p.12). Moreover, according to our interview with an orange exporter, as late arrivals will leave downstream customers with less time to make sales, the exporter has to compensate local distributors for late-arriving fruit. According to Zespri (2014), the company is now using a combination of charter and liner services to deliver its products. As the precise markdown schedule may depend on the agreement between the shipper and downstream distributors, which is not accessible, in the numerical study we assume $p(t) = 100$ for $1 \leq t \leq 10$; $p(t) = 100 - \frac{5}{3}(t - 10)$ for $11 \leq t \leq 70$; $p(N + 1) = p(71) = 0$.

Our analysis in Section 4 has revealed the two benefits of Zespri's current strategy of using a mix of shipping modes. On the one hand, the fruit in chartered vessels arrives early in the season and enjoys a premium, whereas the fruit shipped by ocean liners enables Zespri to sustain the fruit supply till late in the season, which helps improve its overall service level. On the other hand, splitting shipments helps mitigate risks associated with the notoriously unreliable ocean shipping services.

6.2. Impact of Schedule Reliability

To examine how the uncertain arrival schedule influences portfolio selection, we consider the single-demand setting in which $D \sim N(1000, \sigma_D^2)$, where $\sigma_D = 1000cv_D$, $h^O = 0.1$, $h_i^B = 0.05$ and $s_i = 1$ for all $i = 1, 2, \dots, 7$. Recall that in the single-demand setting, it suffices to consider the earliest departure of every shipping service. Thus, $M = 7$. When implementing Algorithm 2, we choose $\hat{T}_i = E[T_i]$ for all i as the initial approximation and $\delta_k = \frac{1}{k+1}$. In the following experiments, we will report the solution generated after 100 iterations of Algorithm 2 as the “optimal” portfolio, denoted by $\mathbf{q}^*$.

Tables 2a and 2b compare the optimal shipping volumes given different levels of cv_T and ρ , where $\mathbf{q}^d$ represents the optimal volumes when $\mathbf{T}$ is perfectly deterministic. It can be seen that compared to $\mathbf{q}^d$, the optimal portfolio can be vastly different when the uncertainty in $\mathbf{T}$ is taken into account. As cv_T increases or ρ decreases, it is imperative to select more services with similar arrival schedules (e.g., SL2, SL3, SL4, SL5) in order to diversify the risks arising from uncertain arrival

times. These observations highlight the importance of incorporating uncertain arrival schedules into the portfolio selection model.

	Charter	SL1	SL2	SL3	SL4	SL5	SL6
$E[T_i]$	24	33	40	42	43	44	46
q^d	668	296	0	0	0	0	552
$q^*(cv_T = .1)$	771	111	0	0	52	165	426
$q^*(cv_T = .2)$	710	0	0	86	152	261	354
$q^*(cv_T = .3)$	578	0	54	135	213	269	358

(a) Impact of cv_T ($\rho = 0$, $cv_D = .6$)

	Charter	SL1	SL2	SL3	SL4	SL5	SL6
$E[T_i]$	24	33	40	42	43	44	46
q^d	668	296	0	0	0	0	552
$q^*(\rho = .8)$	910	0	0	0	15	111	311
$q^*(\rho = .4)$	857	0	0	9	122	125	248
$q^*(\rho = 0)$	814	0	12	30	155	133	229

(b) Impact of ρ ($cv_T = .1$, $cv_D = .4$)

Table 2 Impact of service reliability on the optimal shipment allocation

6.3. Impact of Inventory Holding Costs

Next, we demonstrate the implications of inventory holding costs in the multi-demand setting. All $d(t)$'s are independent and normally distributed where $E[d(1)] = 100$ and $E[d(t)] = 10$ for $t \geq 2$, and the coefficient of variation for each $d(t)$ equals 0.6. The other parameters are unchanged. With multiple demands over time, for each shipping service, we need to consider as candidates all the vessels that depart in different weeks but arrive within the season.

For illustrative purpose, we restrict our attention to the voyages provided by SL3, SL5 and SL6 in Table 1 and set $cv_T = 0.1$ and $\rho = 0$.

The optimal shipping schedule when $h^D = 0.1$																																			
Day ...	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	...
SL 3	228							0							0							0							0						
SL 5			196						0							0							0								0				
SL 6					268						0								0						0									0	

The optimal shipping schedule when $h^D = 0.5$																																			
Day ...	42	43	44	45	46	47	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63	64	65	66	67	68	69	70	71	72	73	74	75	...
SL 3	229							0							0							0								0					
SL 5			123						0							0							0									0			
SL 6					164						36								57						49								0		

Table 3 Impact of h^D on the optimal shipment allocation in the multiple-demand setting

Table 3 presents the optimal portfolios given $h^D = 0.1$ and $h^D = 0.5$. The optimal portfolio under the multi-demand setting can be viewed as a shipping schedule, prescribing the shipping volumes (reported in the table) arriving in each time period. Early in the season all services are

simultaneously used for diversification, and late in the season the shipper may need to defer some shipments with the slow and cheap service (when $h^D = 0.5$) to properly allocate the inventory holding costs incurred at the origin and the destination. However, comparing the cases of $h^D = 0.1$ and $h^D = 0.5$, we find that a change in the value of h^D may lead to different extents of diversification: With a higher h^D , more shipments need to be deferred so as to reduce the inventory cost at the destination.

7. Extensions

7.1. Setup Costs

A setup cost may be incurred when shippers decide to use one more vessel. For example, the setup cost may include the shipper's internal administrative cost involved in signing an additional shipping contract. Moreover, if the vessel has a different departure time than others, the shipper will have to send the cargoes to the port terminal via another truck fleet.

Let K_i denote the setup cost incurred if the shipper uses vessel i , where $i = 1, 2, \dots, M$, and vessels are indexed according to the expected arrival sequence. Suppose that the arrival sequence is deterministic. The cost function can be adapted as

$$H^K(\mathbf{q}) = H(\mathbf{q}) + \sum_{i=1}^M K_i I(q_i > 0) = \sum_{i=1}^M \{L_i(q_i) + c_i q_i + K_i I(q_i > 0)\}, \quad (23)$$

where the L_i 's are defined in Section 5.1.2.

Although the problem with setup costs is nonconvex, the necessary optimality condition stated in Lemma 2 remains valid. Therefore, we can compute the optimal solution as long as the set of vessels to be used has been determined. Moreover, given the optimal portfolio $J = \{j_1, j_2, \dots, j_{|J|}\}$ and the optimal aggregate quantities $\mathbf{y}^*$, the expected cost can be expressed as

$$H^K(\mathbf{y}^*) = (w_{0,j_1} + K_{j_1}) + \sum_{k=1}^{|J|-1} (w_{j_k, j_{k+1}}(y_{j_k}^*) + K_{j_{k+1}}) + w_{j_{|J|}, M+1}(y_{j_{|J|}}^*), \quad (24)$$

where $w_{0,i}$, $w_{i,j}(\cdot)$ and $w_{i,M+1}(\cdot)$ are defined in (17)(18) and (19).

As before, the cost function is decomposed into $|J| + 1$ parts, but now the first $|J|$ parts also include the setup costs incurred by each of the vessels in the portfolio. We can then modify graph $\mathcal{G}$ constructed in Section 5.1.2 and solve (23) in an analogous manner. For the two parameters y_{ij} and d_{ij} of each edge, we only need to adjust the value of d_{ij} to include setup costs while keeping the definition of parameters y_{ij} unchanged, as for any given portfolio J , the resulting aggregate shipping quantities $\mathbf{y}$ also satisfy the necessary optimality conditions in Lemma 2. In particular, we can redefine the parameters d_{ij} as follows:

- for edge (i, j) where $1 \leq i < j \leq M$, $d_{ij} = w_{i,j}(y_{ij}) + K_j$;

- for edge (O, i) where $1 \leq i \leq M$, $d_{O,i} = w_{0,i} + K_i$;
- for edge (i, E) where $1 \leq i \leq M$ and edge (O, E) , the distance parameters are defined as before, i.e., $d_{i,E} = w_{i,M+1}(y_{i,E})$ and $d_{O,E} = \sum_{i=1}^M L_i(0)$.

For an uncertain arrival sequence, the idea of Algorithm 2 can still be borrowed to construct heuristics, although the convergence may not hold due to the nonconvexity of (23). For example, one can initially ignore setup costs and run Algorithm 2 for $\bar{k}$ iterations, where $\bar{k}$ is chosen as a large number so that the resulting approximation $\hat{H}^{\bar{k}}(\mathbf{q})$ is reasonably close to the exact objective function without setup costs. Then, we can obtain an approximate solution by minimizing $\hat{H}^{\bar{k}}(\mathbf{q}) + \sum_{i=1}^M K_i I(q_i > 0)$ over $\mathbf{q} \geq \mathbf{0}$ using the monotone shortest-path algorithm described earlier.

7.2. Single Demand with Partial Lost Sales

In addition to markdowns, late arrivals may also lead to a loss of customers. In this subsection, we consider the single-demand setting ($D(t) = D$ for all t and $h^D = 0$) where some of the unsatisfied demand will be lost. In each time period, we assume that only a fraction β of unsatisfied customers remains in the market. Conditioning on an arrival schedule $\mathbf{T}$ where $T_1 \leq T_2 \leq \dots \leq T_M$, the shipper's revenue from period T_1 to $T_2 - 1$ equals

$$\begin{aligned} & p(T_1) \min\{\beta^{T_1-1} D, q_1\} + p(N+1)(q_1 - \beta^{T_1-1} D)^+ \\ & = (p(T_1) - p(N+1))\beta^{T_1-1} D - (p(T_1) - p(N+1))\beta^{T_1-1} (D - \frac{q_1}{\beta^{T_1-1}})^+ + p(N+1)q_1. \end{aligned}$$

Likewise, for $i = 2, \dots, M$, the revenue from periods T_i to $T_{i+1} - 1$ (where $T_{M+1} = N+1$) can be written as

$$\begin{aligned} & p(T_i)\beta^{T_i-1} \min\{[D - \sum_{k=1}^{i-1} \frac{q_k}{\beta^{T_k-1}}]^+, \frac{q_i}{\beta^{T_i-1}}\} + p(N+1)\beta^{T_i-1} (\frac{q_i}{\beta^{T_i-1}} - [D - \sum_{k=1}^{i-1} \frac{q_k}{\beta^{T_k-1}}]^+)^+ \\ & = (p(T_i) - p(N+1))\beta^{T_i-1} [D - \sum_{k=1}^{i-1} \frac{q_k}{\beta^{T_k-1}}]^+ - (p(T_i) - p(N+1))\beta^{T_i-1} [D - \sum_{k=1}^i \frac{q_k}{\beta^{T_k-1}}]^+ + p(N+1)q_i. \end{aligned}$$

Summing up the above revenue terms and subtracting the shipping cost and inventory holding cost, with partial lost sales, the shipper's problem for fixed $\mathbf{T}$ can be expressed as

$$\begin{aligned} \min_{\mathbf{q} \geq \mathbf{0}} H^L(\mathbf{q}|\mathbf{T}) = & E[\sum_{i=1}^{M-1} ((p(T_i) - p(N+1))\beta^{T_i-1} - (p(T_{i+1}) - p(N+1))\beta^{T_{i+1}-1}) (D - \sum_{k=1}^i \frac{q_k}{\beta^{T_k-1}})^+ \\ & + (p(T_M) - p(N+1))\beta^{T_M-1} (D - \sum_{k=1}^M \frac{q_k}{\beta^{T_k-1}})^+] + \sum_{i=1}^M c_i q_i. \end{aligned} \quad (25)$$

The above problem has the same structure as our single-demand model without lost sales except that the variables and parameters must be adjusted by the fraction β . Specifically, we can define $Q_i = \frac{q_i}{\beta^{T_i-1}}$ as new decision variables and then regard $c_i \beta^{T_i-1}$ and $(p(T_i) - p(N+1))\beta^{T_i-1}$ in (25)

as c_i and $(p(T_i) - p(N + 1))$ in our base model, respectively. Therefore, for a fixed $\mathbf{T}$, Algorithm 1 and Theorem 3 can be extended to incorporate the partial lost sales. To allow for uncertain $\mathbf{T}$, we can follow the idea of Algorithm 2. Conditioning on $\mathbf{T}$, the gradient $\nabla H^L(\mathbf{q}|\mathbf{T})$ is not difficult to evaluate. Thus, one can first construct an initial approximation $\hat{H}^L$ using some fixed nominal arrival schedule, and then calibrate $\mathbf{c}$ iteratively with the gradient information.

It is worth noting that the insights derived from our base model may change when the fraction of lost sales is large. Appendix EC.3 provides a numerical example showing that the value of diversifying with slow services is reduced as the market size rapidly shrinks over time. The important caveat is that although the benefit of using slow services arises from time-sensitive prices, shippers should also pay attention to potential lost sales due to late arrivals.

8. Conclusion

Ocean carriers offer shipping services with various transit times and freight rates. A major challenge facing seasonal product exporters is how to balance the trade-off between the speed and the transportation cost. Delivering the product to the market earlier normally brings higher profit margins, but faster shipping services are more costly. Given uncertainties in both demand and vessels' arrival times, we show that maintaining a portfolio of shipping services enables the shipper to manage risks from both the demand and supply sides. On the one hand, a combination of fast and slow shipping modes helps the shipper cope with demand uncertainty in a similar fashion that a portfolio of supply option contracts does in the newsvendor procurement problem (e.g., Martínez-de Albéniz and Simchi-Levi 2005). On the other hand, splitting shipments among vessels also serves to mitigate the uncertainty in arrival schedules, which mirrors the effect of supply diversification (e.g., Anupindi and Akella 1993). Therefore, our shipping problem combines these two portfolio effects which have been previously understood as separate issues. By extending several existing results in the literature, we have first analyzed the structure of the solution under some conditions, and then developed a solution approach to the general problem. Moreover, a real-world case study has been conducted to demonstrate the practical implications of our model.

There are several directions for future research. First, we have assumed an exogenous demand process with a given price schedule. In reality, a powerful seasonal-product seller may be able to influence the market price. Endogenizing pricing decisions will be an interesting extension for future study. Second, another reason for shippers to use multiple vessels could be the limited capacity of ocean carriers. The problem becomes more challenging if capacity constraints are taken into account. Third, we have assumed that the markdown process is exogenously given. The case in which the shipper can also determine market prices is an interesting direction for future research. Finally, we may further endogenize carriers' decisions based on a game-theoretical framework, as

the shipper would be able to obtain lower freight rates and better services from the competition among carriers.

Endnotes

1. The contract between a shipper and a carrier often only specifies a price per TEU (twenty-foot equivalent unit). The shipper may also incur an internal fixed cost when shipping via an additional vessel. We will incorporate this as the setup cost in Section 7.1.
2. Throughout the paper, we use “increasing/decreasing” in a weak sense.
3. Model III in Anupindi and Akella (1993) considers a multi-period dual sourcing problem in which the delivery can be delayed. However, the length of delay can only be one period.
4. For example, Babich et al. (2007) present a similar graphic illustration for their Proposition 1.
5. For instance, one may set $\bar{Q}_i$ as the newsvendor solution when vessel i is the only candidate. Note that $H(\mathbf{q})$ is supermodular in $\mathbf{q}$ and therefore q_i^* cannot exceed the optimal quantity in vessel i when all other vessels are inactive.
6. http://www.tis-gdv.de/tis_e/ware/obst/kiwi/kiwi.htm
7. <http://www.sea-distances.org/>. Charter services may also make intermediate stops. Here, we assume an exclusive charter service, sailing directly from Auckland to Rotterdam, to represent the fastest service.

Acknowledgments

We are very grateful to Professor Chung-Piaw Teo (Area Editor), Associate Editor and anonymous referees’ constructive comments which have helped improve this paper significantly. The work described in this paper was fully supported by a grant from the Research Grants Council of the HKSAR, China, T32-620/11. The corresponding author of this paper is Prof. Chung-Yee Lee.

References

- Agricultural Marketing Service (2012) Fruit and vegetable market news user guid. Available at <http://www.ams.usda.gov/AMSv1.0/getfile?dDocName=STELPRDC5048679>.
- Anupindi R, Akella R (1993) Diversification under supply uncertainty. *Management Science* 39(8):944–963.
- Babich V, Burnetas AN, Ritchken PH (2007) Competition and diversification effects in supply chains with supplier default risk. *Manufacturing & Service Operations Management* 9(2):123–146.
- Butler Z (2012) Single source shortest path: The bellman-ford algorithm (lecture notes). Available at <http://www.cs.rit.edu/~zjb/courses/800/lec15-1.pdf>.
- Cheung RKM, Powell WB (2000) Shape-a stochastic hybrid approximation procedure for two-stage stochastic programs. *Operations Research* 48(1):73–79.

- Dada M, Petruzzi NC, Schwarz LB (2007) A newsvendors procurement problem when suppliers are unreliable. *Manufacturing & Service Operations Management* 9(1):9–32.
- Federgruen A, Yang N (2008) Selecting a portfolio of suppliers under demand and supply risks. *Operations Research* 56(4):916–936.
- Federgruen A, Yang N (2009) Optimal supply diversification under general supply risks. *Operations Research* 57(6):1451–1468.
- Feng Q, Sethi S, Yan H, Zhang H (2006) Are base-stock policies optimal in inventory problems with multiple delivery modes? *Operations Research* 54(4):801–807.
- Fu Q, Lee CY, Teo CP (2010) Procurement management using option contracts: Random spot price and the portfolio effect. *IIE transactions* 42(11):793–811.
- Fukuda Y (1964) Optimal policies for the inventory problem with negotiable leadtime. *Management Science* 10(4):690–708.
- Kaplan RS (1970) A dynamic inventory model with stochastic lead times. *Management Science* 16(7):491–507.
- Kouvelis P, Li J (2008) Flexible backup supply and the management of lead-time uncertainty. *Production and Operations Management* 17(2):184–199.
- Lau HS, Zhao LG (1993) Optimal ordering policies with two suppliers when lead times and demands are all stochastic. *European Journal of Operational Research* 68(1):120–133.
- Lavigne GM (2014) Cargosmart: July-september vessel delays at manila averaged five days. *Journal of Commerce* Available at http://www.joc.com/port-news/asian-ports/cargosmart-july-september-vessel-delays-manila-averaged-five-days_20141019.html.
- Lim A, Rodrigues B, Xu Z (2008) Transportation procurement with seasonally varying shipper demand and volume guarantees. *Operations research* 56(3):758–771.
- Martínez-de Albéniz V, Simchi-Levi D (2005) A portfolio approach to procurement contracts. *Production and Operations Management* 14(1):90–114.
- Martínez-de Albéniz V, Simchi-Levi D (2009) Competition in the supply option market. *Operations Research* 57(5):1082–1097.
- Mechelen Y (2011) Reefers come out on top when it comes to kiwi fruit. Available at <http://www.tradewindsnews.com/weekly/149672/reefers-come-out-on-top-when-it-comes-to-kiwi-fruit>.
- Ramasesh RV, Ord JK, Hayya JC (1993) Note: dual sourcing with nonidentical suppliers. *Naval Research Logistics* 40(2):279–288.
- Ramasesh RV, Ord JK, Hayya JC, Pan A (1991) Sole versus dual sourcing in stochastic lead-time (s, q) inventory models. *Management science* 37(4):428–443.

- Schultz A (2011) Congruence; sylvesters law of inertia (lecture notes). Available at http://palmer.wellesley.edu/~aschultz/w10/math416/coursenotes_and_handouts/coursenotes_100422.pdf.
- Sedarage D, Fujiwara O, Luong HT (1999) Determining optimal order splitting and reorder level for n - supplier inventory systems. *European Journal of Operational Research* 116(2):389–404.
- Sharypova K, van Woensel T, Fransoo JC (2012) Coordination and analysis of barge container hinterland networks, working paper, Eindhoven University of Technology, Netherlands.
- Snyder LV, Atan Z, Peng P, Rong Y, Schmitt AJ, Sinsoysal B (2010) Or/ms models for supply chain disruptions: A review, working paper. Lehigh University.
- Tomlin B (2006) On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* 52(5):639–657.
- Tomlin B, Wang Y (2010) Operational strategies for managing supply chain disruption risk. P. Kouvelis, O. Boyabatli, L. Dong, R. Li, eds, *The Handbook of Integrated Risk Management in Global Supply Chains* (John Wiley & Sons, Inc, New York).
- Van de Weijer FA (2013) Improving supply chain cost and hinterland modal split: An inventory theoretic approach, master Thesis, Eindhoven University of Technology, Netherlands.
- Wan K, Levary RR (1995) A linear-programming-based price negotiation procedure for contracting shipping companies. *Transportation Research Part A: Policy and Practice* 29(3):173–186.
- Whittimore A, Saunders S (1977) Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics* 32(2):293–305.
- Zespri (2014) Zespri annual review 2013/14. Available at <https://www.zespri.com/ZespriInvestorPublications/annual-review-2013-14.pdf>.
- Zipkin P (1986) Stochastic leadtimes in continuous-time inventory models. *Naval Research Logistics Quarterly* 33(4):763–774.

Tao Lu is an assistant professor at Rotterdam School of Management at Erasmus University. His research interests are in supply chain management, ocean transport logistics, and operations in the sharing economy.

Jan C. Fransoo is a Professor of Operations Management and Logistics in the School of Industrial Engineering at Eindhoven University of Technology in the Netherlands. His research addresses a variety of topics in supply chain management, with a current focus on container supply chains, emerging market logistics, and sustainability.

Prof. Chung-Yee Lee is currently a Chair Professor/Cheong Ying Chan Professor of Engineering with the Department of Industrial Engineering and Logistics Management, Hong Kong University of Science and Technology (HKUST). He is also Senior Advisor to the President and

Director of the Office of Institutional Research at HKUST. His search areas are in logistics and supply chain management, scheduling and inventory management, with a current focus on ocean transport logistics and e-commerce logistics.

Electronic Companion to “Carrier Portfolio Management for Shipping Seasonal Products”

EC.1. Proofs

Proof of Theorem 1 The theorem is a special case of Theorem 2 when $T_1 \leq T_2$ with probability one. The proof of Theorem 2 will be presented later. Here we show that Theorem 2 reduces to Theorem 1 if $\Pr\{T_1 \leq T_2\} = 1$.

First, notice that for all i , we have

$$E[p(T_i)] - p(N+1) = \sum_{t=1}^N b(t)G_i(t) \quad (\text{EC.1})$$

Invoking equation (EC.1) in part (i) of Theorem 2 yields $\mathbf{q}^* = (0, 0)$ if $c_i \geq E[p(T_i)] - p(N+1)$. From part (ii) of Theorem 2, $\mathbf{q}^* = (q_1^s, 0)$ if $c_1 < \sum_{t=1}^N b(t)G_1(t) = E[p(T_1)] - p(N+1)$ and

$$\begin{aligned} & c_2 - \frac{\sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_1(t)} c_1 - \sum_{t=1}^N b(t)Pr\{T_1 > t, T_2 \leq t\} \\ &= c_2 - \frac{\sum_{t=1}^N b(t)Pr\{T_2 \leq t\}}{\sum_{t=1}^N b(t)G_1(t)} c_1 - 0 \\ &= c_2 - \frac{E[p(T_2)] - p(N+1)}{E[p(T_1)] - p(N+1)} c_1 \geq 0, \end{aligned}$$

where the first equality follows by the assumption of $Pr\{T_1 \leq T_2\} = 1$, and the second equality follows from equation (EC.1). Rearranging the terms, we have $\mathbf{q}^* = (q_1^s, 0)$ if $c_1 < E[p(T_1)] - p(N+1)$ and $\frac{c_1}{E[p(T_1)] - p(N+1)} \leq \frac{c_2}{E[p(T_2)] - p(N+1)}$.

Similarly, from part (ii) of Theorem 2, $\mathbf{q}^* = (0, q_2^s)$ if $c_2 < E[p(T_2)] - p(N+1)$ and

$$\begin{aligned} & c_1 - \frac{\sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_2(t)} c_2 - \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 > t\} \\ &= c_1 - c_2 - \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 > t\} \\ &= c_1 - c_2 - \sum_{t=1}^N b(t)(Pr\{T_1 \leq t\} - Pr\{T_1 \leq t, T_2 \leq t\}) \\ &= c_1 - c_2 - \sum_{t=1}^N b(t)(Pr\{T_1 \leq t\} - Pr\{T_2 \leq t\}) \\ &= c_1 - c_2 - E[p(T_1)] + E[p(T_2)] \geq 0, \end{aligned}$$

where we have also invoked the assumption $Pr\{T_1 \leq T_2\} = 1$ in the derivation. Hence, $\mathbf{q}^* = (0, q_2^s)$ if $c_2 < E[p(T_2)] - p(N+1)$ and $E[p(T_1)] - c_1 \leq E[p(T_2)] - c_2$.

Finally, from part(iii) of Theorem 2, we have $\mathbf{q}^* = (q_1^d, q_2^d)$ if $E[p(T_1)] - c_1 > E[p(T_2)] - c_2$ and $\frac{c_1}{E[p(T_1)] - p(N+1)} > \frac{c_2}{E[p(T_2)] - p(N+1)}$, i.e., conditions (5) and (6) hold.

We may also check whether the two equations, $q_1^d = \bar{F}^{-1}(\frac{c_1 - c_2}{E[p(T_1)] - E[p(T_2)]})$ and $q_1^d + q_2^d = \bar{F}^{-1}(\frac{c_2}{E[p(T_2)] - p(N+1)})$, have a solution. It suffices to verify that $\frac{c_2}{E[p(T_2)] - p(N+1)} < \frac{c_1 - c_2}{E[p(T_1)] - E[p(T_2)]}$. We have

$$\begin{aligned} \frac{c_2}{E[p(T_2)] - p(N+1)} &< \frac{c_1 - c_2}{E[p(T_1)] - E[p(T_2)]} \\ \Leftrightarrow c_2(E[p(T_1)] - E[p(T_2)]) &< (c_1 - c_2)(E[p(T_2)] - p(N+1)) \\ \Leftrightarrow c_2 E[p(T_1)] - c_2 E[p(T_2)] &< c_1 E[p(T_2)] - c_1 p(N+1) - c_2 E[p(T_2)] + c_2 p(N+1) \\ \Leftrightarrow c_2(E[p(T_1)] - p(N+1)) &< c_1(E[p(T_2)] - p(N+1)) \\ \Leftrightarrow \frac{c_2}{E[p(T_2)] - p(N+1)} &< \frac{c_1}{E[p(T_1)] - p(N+1)}. \end{aligned}$$

Hence, the existence of solutions is guaranteed by condition (6).

Proof of Corollary 1 Part (i) follows immediately from Theorem 1.

Let us consider part (ii). By mean-preserving transform, we can write D as $D = \eta \hat{D} + (1 - \eta)\mu_D$ where $0 \leq \eta \leq 1$. Note that $\text{Var}(D) = \eta^2 \text{Var}(\hat{D})$ so $\text{Var}(D)$ is increasing in η . Denote by $F_D(x)$ and $F_{\hat{D}}(x)$ the cumulative functions of D and $\hat{D}$, respectively. By Theorem 1, in the optimal solution we have $\bar{F}_D(q_1^d) = \bar{F}_{\hat{D}}(\hat{q}_1^d)$ and $\bar{F}_D(q_1^d + q_2^d) = \bar{F}_{\hat{D}}(\hat{q}_1^d + \hat{q}_2^d)$ because no parameters change except D . Notice that

$$\begin{aligned} \bar{F}_D(x) &= \Pr\{D > x\} = \Pr\{\eta \hat{D} + (1 - \eta)\mu_D > x\} \\ &= \Pr\{\hat{D} > \frac{x - (1 - \eta)\mu_D}{\eta}\} \\ &= \bar{F}_{\hat{D}}(\frac{x - (1 - \eta)\mu_D}{\eta}). \end{aligned}$$

Hence, $\hat{q}_1^d = \frac{q_1^d - (1 - \eta)\mu_D}{\eta}$ and $\hat{q}_1^d + \hat{q}_2^d = \frac{q_1^d + q_2^d - (1 - \eta)\mu_D}{\eta}$. These equations together imply $q_2^d = \eta \hat{q}_2^d \leq \hat{q}_2^d$. Also, by rearranging the terms, we have $q_1^d - \mu_D = \eta(\hat{q}_1^d - \mu_D)$ and $q_1^d + q_2^d - \mu_D = \eta(\hat{q}_1^d + \hat{q}_2^d - \mu_D)$. Then, the results immediately follow.

For part (iii), without loss of generality, suppose that the transit time of vessel 1 is shorter, i.e., $t_0 + E[T_1] - 1 - s_1 \leq t_0 + E[T_2] - 1 - s_2$, or equivalently, $E[T_1] - s_1 \leq E[T_2] - s_2$. Then, $\frac{c_1 - c_2}{E[p(T_1)] - E[p(T_2)]}$ is decreasing in h^B , as $\frac{d}{dh^B}(c_1 - c_2) = (E[T_1] - s_1) - (E[T_2] - s_2) \leq 0$. On the other hand, $\bar{F}^{-1}(x)$ is decreasing in x . Hence, q_1^d is increasing in h^B . Similarly, since $q_1^d + q_2^d = \bar{F}^{-1}(\frac{c_2}{E[p(T_2)] - p(N+1)})$, and c_2 is increasing in h^B , the total quantity $q_1^d + q_2^d$ is decreasing in h^B . It then follows that q_2^d is decreasing in h^B .

Proof of Theorem 2 Let (q_1^*, q_2^*) denote the optimal solution to (2). It can be verified that $H(\mathbf{q})$ is convex in $\mathbf{q}$. Thus, the first-order condition is necessary and sufficient for the optimality. Then, it suffices to verify that $\frac{\partial H(\mathbf{q}^*)}{\partial q_i} = 0$ if $q_i^* > 0$ and $\frac{\partial H(\mathbf{q}^*)}{\partial q_i} \geq 0$ if $q_i^* = 0$ for $i = 1, 2$.

Taking the first derivative, we have, for $i = 1, 2$,

$$\frac{\partial H(\mathbf{q})}{\partial q_i} = c_i + \sum_{t=1}^N b(t) \frac{\partial \mathbb{E}[D - \sum_{i=1}^2 q_i I(T_i \leq t)]^+}{\partial q_i}, \quad (\text{EC.2})$$

where

$$\frac{\partial \mathbb{E}[D - \sum_{i=1}^2 q_i I(T_i \leq t)]^+}{\partial q_i} = -\bar{F}(q_1 + q_2) \Pr\{T_i \leq t, T_{3-i} \leq t\} - \bar{F}(q_i) \Pr\{T_i \leq t, T_{3-i} > t\}.$$

All possible pairs of (c_1, c_2) can be segmented into four regions according to the conditions in cases (i)-(iii) (see Figure 4). In what follows, we will verify the first-order condition with the proposed solution for each segmentation.

Under the conditions in case (i), we have $\frac{\partial H(\mathbf{q})}{\partial q_i} = c_i - \sum_t b(t)G_i(t) \geq 0$ for all $i = 1, 2$ when $\mathbf{q} = (0, 0)$. The first-order condition is satisfied at $\mathbf{q} = (0, 0)$ and this implies that $\mathbf{q}^* = (0, 0)$ is the optimal solution in case (i). Furthermore, $\frac{\partial \mathbb{E}[D - \sum_{i=1}^2 q_i I(T_i \leq t)]^+}{\partial q_i}$ is increasing in both q_1 and q_2 and so is $\frac{\partial H(\mathbf{q}^*)}{\partial q_i}$. Therefore, if $q_i > 0$ for some i , the first-order derivative $\frac{\partial H(\mathbf{q}^*)}{\partial q_i}$ cannot equal zero. Hence, $\mathbf{q}^* = (0, 0)$ is also the unique solution in case (i).

For case (ii), we consider $i = 1$ for instance. For $i = 2$, the proof is symmetric. We need to verify that $\frac{\partial H(\mathbf{q})}{\partial q_1} = 0$ and $\frac{\partial H(\mathbf{q})}{\partial q_2} \geq 0$ when $q_1 = q_1^s$ and $q_2 = 0$. Given $q_1 = q_1^s$ and $q_2 = 0$, it follows that

$$\frac{\partial H(\mathbf{q})}{\partial q_1} = c_1 - \sum_t b(t)G_1(t)\bar{F}(q_1^s) = 0,$$

where the equality follows immediately by (7). Also, we have

$$\begin{aligned} \frac{\partial H(\mathbf{q})}{\partial q_2} &= c_2 - \bar{F}(q_1^s) \sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 \leq t\} - \bar{F}(0) \sum_{t=1}^N b(t) \Pr\{T_1 > t, T_2 \leq t\} \\ &= c_2 - c_1 \frac{\sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_1(t)} - \sum_{t=1}^N b(t) \Pr\{T_1 > t, T_2 \leq t\} \\ &\geq 0 \end{aligned}$$

where the third equality follows by invoking (7) and $\bar{F}(0) = 1$, and the last inequality follows provided that the condition of case (ii) holds. By convexity, we can conclude that $(q_1^s, 0)$ is optimal in case (ii) when $i = 1$. Furthermore, we can also argue that $(q_1^s, 0)$ is the unique solution. First, $(0, 0)$ cannot be optimal in this case as $\frac{\partial H(\mathbf{q})}{\partial q_1}|_{\mathbf{q}=(0,0)} = c_1 - \sum_{t=1}^N b(t)G_1(t) < 0$, which contradicts the optimality condition. Second, suppose $q_1 = 0$ and $q_2 > 0$ in the optimal solution. We must have $\frac{\partial H(\mathbf{q})}{\partial q_2} = 0$, which implies $\bar{F}(q_2) = \frac{c_2}{\sum_{t=1}^N b(t)G_2(t)} \leq 1$. The optimality also requires

$$\begin{aligned} \frac{\partial H(\mathbf{q})}{\partial q_1}|_{\mathbf{q}=(0,q_2)} &= c_1 - \sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 \leq t\} \bar{F}(q_2) - \sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 > t\} \\ &= c_1 - \frac{\sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_2(t)} c_2 - \sum_{t=1}^N b(t) \Pr\{T_1 \leq t, T_2 > t\} \geq 0. \end{aligned}$$

However, any pair of (c_1, c_2) satisfying the above inequality would not fall into the set $\{(c_1, c_2) : c_1 - \sum_{t=1}^N b(t)G_1(t) < 0, c_2 - c_1 \frac{\sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_1(t)} - \sum_{t=1}^N b(t)Pr\{T_1 > t, T_2 \leq t\}\}$. This can be checked easily by drawing a graph in the c_1 - c_2 coordinate plane. Finally, suppose $q_1 > 0$ and $q_2 > 0$ in the optimal solution. We must have $c_i - \bar{F}(q_1 + q_2) \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\} - \bar{F}(q_i) \sum_{t=1}^N b(t)Pr\{T_i \leq t, T_{3-i} > t\} = 0$ for both $i = 1, 2$. The equation for $i = 1$ implies that

$$\begin{aligned} 0 &= c_1 - \bar{F}(q_1 + q_2) \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\} - \bar{F}(q_1) \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 > t\} \\ &< c_1 - \bar{F}(q_1 + q_2) \sum_{t=1}^N b(t)G_1(t). \end{aligned}$$

Equivalently, $\bar{F}(q_1 + q_2) < \frac{c_1}{\sum_{t=1}^N b(t)G_1(t)}$. The equation for $i = 2$ requires

$$\begin{aligned} 0 &= c_2 - \bar{F}(q_1 + q_2) \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\} - \bar{F}(q_2) \sum_{t=1}^N b(t)Pr\{T_2 \leq t, T_1 > t\} \\ &> c_2 - \bar{F}(q_1 + q_2) \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\} - \sum_{t=1}^N b(t)Pr\{T_2 \leq t, T_1 > t\} \\ &> c_2 - \frac{c_1 \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_1(t)} - \sum_{t=1}^N b(t)Pr\{T_2 \leq t, T_1 > t\}, \end{aligned}$$

where the first inequality follows as $\bar{F}(q_2) < 1$ for $q_2 > 0$ and the second inequality follows as $\bar{F}(q_1 + q_2) < \frac{c_1}{\sum_{t=1}^N b(t)G_1(t)}$. However, the resulting inequality contradicts the condition in case (ii) when $i = 1$. Therefore, $q_1 > 0$ and $q_2 > 0$ cannot be optimal in case (ii). In sum, $(q_1^*, 0)$ is also the unique optimal solution in case (ii) when $i = 1$.

In case (iii), we have $c_{3-i} - \frac{c_i \sum_{t=1}^N b(t)Pr\{T_1 \leq t, T_2 \leq t\}}{\sum_{t=1}^N b(t)G_i(t)} - \sum_{t=1}^N b(t)Pr\{T_{3-i} \leq t, T_i > t\} < 0$ for both $i = 1, 2$ (corresponding to region *ABOE* in Figure 4). These two conditions in effect guarantee that there exist $q_1^* > 0$ and $q_2^* > 0$ such that $\frac{\partial H(\mathbf{q}^*)}{\partial q_i} = 0$ for both $i = 1, 2$. Clearly, $q_i = 0$ for some i cannot be optimal in case (iii), since the optimality of the boundary solution requires $\frac{\partial H(\mathbf{q}^*)}{\partial q_i} \geq 0$ and $q_i^* = 0$, which can be achieved only under the conditions in case (i) or (ii).

Proof of Theorem 3 Due to the convexity of the objective function, it suffices to verify the first-order conditions, which are given by $\mathbf{q}^{*T} \nabla H(\mathbf{q}^*) = 0$, i.e., $\frac{\partial H(\mathbf{q}^*)}{\partial q_i^*} = 0$ for $q_i^* > 0$ and $\frac{\partial H(\mathbf{q}^*)}{\partial q_i} \geq 0$ for $q_i^* = 0$.

The Case of $J = \{i_A\}$ First, consider the case of $|J| = 1$, i.e., $i_A = i_B$. The solution to be verified is $q_{i_A}^* = \bar{F}^{-1}(\frac{c_{i_A}}{\bar{p}_{i_A} - p(N+1)})$, and $q_i^* = 0$ for any $i \neq i_A$. So it suffices to show that $\frac{\partial H}{\partial q_i} = 0$ for $i = i_A$ and $\frac{\partial H}{\partial q_i} \geq 0$ for any $i \neq i_A$.

For any $i > i_A$, the first-order derivative is given by

$$\begin{aligned}\frac{\partial H}{\partial q_i} \Big|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - p(N+1))\bar{F}(q_1^* + q_2^* + \dots + q_i^*) \\ &= c_i - (\bar{p}_i - p(N+1))\bar{F}(q_{i_A}^*) \geq 0\end{aligned}$$

where the inequality follows by (11), the definition of i_B (recall that in this case i_A and i_B denote the same shipping service).

For $i = i_A$, we have

$$\frac{\partial H}{\partial q_{i_A}} \Big|_{\mathbf{q}=\mathbf{q}^*} = c_{i_A} - (\bar{p}_{i_A} - p(N+1))\bar{F}(q_{i_A}^*) = 0.$$

For $i < i_A$, we have

$$\begin{aligned}\frac{\partial H}{\partial q_i} \Big|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - \bar{p}_{i_A})\bar{F}(0) - (\bar{p}_{i_A} - p(N+1))\bar{F}(q_{i_A}^*) \\ &= c_i - (\bar{p}_i - \bar{p}_{i_A}) - c_{i_A} \geq 0,\end{aligned}$$

where the last inequality holds by (10), the definition of i_A . Hence, we have completed the proof for the case of $i_A = i_B$.

The Case of $J = \{i_A, i_B\}$ Consider the case of $J = \{i_A, i_B\}$. Similarly, we need to show $\frac{\partial H}{\partial q_i} \geq 0$ for any $i \neq i_A$ and i_B , and $\frac{\partial H}{\partial q_i} = 0$ for $i = i_A$ and $i = i_B$.

For $i > i_B$, the first-order derivative is

$$\begin{aligned}\frac{\partial H}{\partial q_i} \Big|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - p(N+1))\bar{F}(q_1^* + q_2^* + \dots + q_i^*) \\ &= c_i - (\bar{p}_i - p(N+1))\bar{F}(q_{i_A}^* + q_{i_B}^*) \\ &= c_i - (\bar{p}_i - p(N+1))\frac{c_{(i_B)}}{\bar{p}_{i_B} - p(N+1)} \geq 0,\end{aligned}$$

where the inequality follows by (11).

For $i = i_B$, we have $\frac{\partial H}{\partial q_{(i_B)}} \Big|_{\mathbf{q}=\mathbf{q}^*} = c_{i_B} - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + q_{i_B}^*) = 0$.

For $i_A < i < i_B$, we have

$$\begin{aligned}\frac{\partial H}{\partial q_i} \Big|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - \bar{p}_{i_B})\bar{F}(q_{i_A}^*) - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + q_{i_B}^*) \\ &= c_i - (\bar{p}_i - \bar{p}_{i_B})\frac{c_{i_A} - c_{i_B}}{\bar{p}_{i_A} - \bar{p}_{i_B}} - (\bar{p}_{i_B} - p(N+1))\frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \\ &= (c_i - c_{i_B}) - (\bar{p}_i - \bar{p}_{i_B})\frac{c_{i_A} - c_{i_B}}{\bar{p}_{i_A} - \bar{p}_{i_B}} \geq 0.\end{aligned}$$

The last inequality can be explained as follows. $J = \{i_A, i_B\}$ implies that Algorithm 1 stops at Step 1. According to Step 1, we must have $\frac{c_{i_A} - c_{i_B}}{\bar{p}_{i_A} - \bar{p}_{i_B}} \leq \frac{c_i - c_{i_B}}{\bar{p}_i - \bar{p}_{i_B}}$ for all $i_A < i < i_B$, which implies the above " ≥ 0 ".

For $i = i_A$, one can directly verify that $\frac{\partial H}{\partial q_{i_A}} \Big|_{\mathbf{q}=\mathbf{q}^*} = c_{i_A} - (\bar{p}_{i_A} - \bar{p}_{i_B})\bar{F}(q_{i_A}^*) - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + q_{i_B}^*) = 0$.

For $i < i_A$, we have

$$\begin{aligned} \left. \frac{\partial H}{\partial q_i} \right|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - \bar{p}_{i_A})\bar{F}(0) - (\bar{p}_{i_A} - \bar{p}_{i_B})\bar{F}(q_{i_A}^*) - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + q_{i_B}^*) \\ &= c_i - (\bar{p}_i - \bar{p}_{i_A}) - c_{i_A} = (\bar{p}_{i_A} - c_{i_A}) - (\bar{p}_i - c_i) \geq 0, \end{aligned}$$

where the last inequality follows by (10), the definition of i_A (service i_A has the highest marginal profit). We are done for the case of $J = \{i_A, i_B\}$.

The Case of $J = \{i_A, j_k, \dots, j_1, i_B\}$, ($k \geq 1$) Similarly, we need to show $\frac{\partial H}{\partial q_i} \geq 0$ for any $i \notin J$, and $\frac{\partial H}{\partial q_i} = 0$ for any $i \in J$.

For $i \geq i_B$, the proof is the same as that for the case of $J = \{i_A, i_B\}$ for $i \geq i_B$.

For $j_1 < i < i_B$, we have

$$\begin{aligned} \left. \frac{\partial H}{\partial q_i} \right|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - \bar{p}_{i_B})\bar{F}(q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^*) - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^* + q_{i_B}^*) \\ &= c_i - (\bar{p}_i - \bar{p}_{i_B}) \frac{c_{j_1} - c_{i_B}}{\bar{p}_{j_1} - \bar{p}_{i_B}} - (\bar{p}_{i_B} - p(N+1)) \frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \\ &= (c_i - c_{i_B}) - (\bar{p}_i - \bar{p}_{i_B}) \frac{c_{j_1} - c_{i_B}}{\bar{p}_{j_1} - \bar{p}_{i_B}} \geq 0, \end{aligned}$$

where the last inequality holds since according to Algorithm 1 j_1 is selected such that $\frac{c_{j_1} - c_{i_B}}{\bar{p}_{j_1} - \bar{p}_{i_B}} \leq \frac{c_i - c_{i_B}}{\bar{p}_i - \bar{p}_{i_B}}$ for all $i_A \leq i < i_B$.

For $j_\tau < i < j_{\tau-1}$ for $\tau = 2, \dots, k$, we have

$$\begin{aligned} \left. \frac{\partial H}{\partial q_i} \right|_{\mathbf{q}=\mathbf{q}^*} &= c_i - (\bar{p}_i - \bar{p}_{j_{\tau-1}})\bar{F}(q_{i_A}^* + \sum_{i=j_\tau}^{j_k} q_i^*) \\ &\quad - \dots - (\bar{p}_{j_1} - \bar{p}_{i_B})\bar{F}(q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^*) - (\bar{p}_{i_B} - p(N+1))\bar{F}(q_{i_A}^* + \sum_{i=j_1}^{j_k} q_i^* + q_{i_B}^*) \\ &= c_i - (\bar{p}_i - \bar{p}_{j_{\tau-1}})\bar{F}(q_{i_A}^* + \sum_{i=j_\tau}^{j_k} q_i^*) - \dots - (c_{j_1} - c_{i_B}) - c_{i_B} \\ &= (c_i - c_{j_{\tau-1}}) - (\bar{p}_i - \bar{p}_{j_{\tau-1}})\bar{F}(q_{i_A}^* + \sum_{i=j_\tau}^{j_k} q_i^*) \\ &= (c_i - c_{j_{\tau-1}}) - (\bar{p}_i - \bar{p}_{j_{\tau-1}}) \frac{c_{j_\tau} - c_{j_{\tau-1}}}{\bar{p}_{j_\tau} - \bar{p}_{j_{\tau-1}}} \geq 0, \end{aligned}$$

where the inequality follows since j_τ is selected such that $\frac{c_{j_\tau} - c_{j_{\tau-1}}}{\bar{p}_{j_\tau} - \bar{p}_{j_{\tau-1}}} \leq \frac{c_i - c_{j_{\tau-1}}}{\bar{p}_i - \bar{p}_{j_{\tau-1}}}$ for all $i_A \leq j_{\tau-1}$.

For $i = j_\tau$ where $\tau = 1, \dots, k$, it is straightforward to verify that $\frac{\partial H}{\partial q_i} = 0$.

For $i_A < i < j_k$, the proof is analogous to that for the case of $J = \{i_A, i_B\}$ for any $i_A < i < i_B$.

For $i \leq i_A$, the proof is analogous to that for the case of $J = \{i_A, i_B\}$ for $i \leq i_A$.

Proof of Corollary 2 From (11), the definition of i_B , one can deduce that

$$\frac{c_{i_B}}{\bar{p}_{i_B} - p(N+1)} \leq \frac{c_{j_1}}{\bar{p}_{j_1} - p(N+1)} \Rightarrow c_{j_1} \geq c_{i_B} \frac{\bar{p}_{j_1} - p(N+1)}{\bar{p}_{i_B} - p(N+1)} \geq c_{i_B}.$$

Likewise, from the constructions of $j_1, j_2, \dots, j_k$, one can verify that $c_{i_A} \geq c_{j_k} \geq \dots \geq c_{j_2} \geq c_{j_1}$.

Proof of Lemma 2 The derivative of $L_i(y)$ is given by $L'_i(y) = -E_{\mathbf{T}}[\sum_{t=T(i)}^{T(i+1)-1} b(t)\bar{F}_t(y) - h^D F_t(y)]$, where we have interchanged the derivative and expectation as the derivative inside the expectation is clearly finite. Given the optimal portfolio $|J|$, we have $\frac{\partial H(\mathbf{q}^*)}{\partial q_{jv}} = 0$ for all $v = 1, 2, \dots, |J|$. These $|J|$ equations constitute a linear system with respect to $\sum_{i=j_k}^{j_{k+1}-1} L'_i(\sum_{v=1}^k q_{jv}^*)$ where $k = 1, 2, \dots, |J| - 1$ and $\sum_{i=j_{|J|}}^M L'_i(\sum_{v=1}^{|J|} q_{jv}^*)$. Equations (14) and (15) immediately follow by solving this linear system.

Proof of Theorem 5 Algorithm 2 falls into a class of stochastic approximation schemes proposed in Cheung and Powell (2000). The only modification is that we use a sample average to estimate the actual gradient, instead of a single sample. This is because in our problem, compared to solving an approximate problem in one iteration, it is relatively simple to evaluate the gradient as the closed-form expression is available. Their proof for the convergence remains valid under this slight modification, and thus we will not repeat it. According to Theorem 1 in their work, the convergence is guaranteed if (1) the constraint set $\mathcal{Q} = \{\mathbf{0} \leq \mathbf{q} \leq \bar{\mathbf{Q}}\}$ is convex and compact; (2) $E[H(\mathbf{q})]$ is convex, finite and continuous on $\mathcal{Q}$; (3) $\nabla H(\mathbf{q}|\mathbf{T})$ is bounded for all realizations of $\mathbf{T}$; (4) $\hat{H}^k(\mathbf{q})$ is strongly convex; (5) the step size satisfies $0 < \delta_k < 1$, $\sum_{k=1}^{\infty} \delta_k = \infty$, $\sum_{k=1}^{\infty} \delta_k^2 < \infty$; (6) $\hat{H}^0(\mathbf{q})$ is bounded and continuous, and $\nabla \hat{H}^0(\mathbf{q})$ is bounded for all $\mathbf{q} \in \mathcal{Q}$. It is straightforward to verify that the gradient of $H(\mathbf{q}|\mathbf{T})$ is bounded for any realization of $\mathbf{T}$ and so is the gradient of $\hat{H}^0$. The only remaining task is to provide conditions for the strong convexity of $\hat{H}^0$.

A function f is strongly convex if and only if there exists a constant $\epsilon > 0$ such that $f(\mathbf{x}) - \epsilon \|\mathbf{x}\|^2$. For a twice-differentiable function, this means that f is strongly convex if $\nabla^2 f(x) \succeq \epsilon I$ for some constant $\epsilon > 0$. Because a strongly convex function is still strongly convex after adding linear terms, it suffices to show that the initial approximation $\hat{H}^0(\mathbf{q})$ is strongly convex. Define $a_k = \sum_{t=\hat{T}_k}^{\hat{T}_{k+1}-1} (b(t) + h^D) f_t(\sum_{j=1}^k q_j)$. For notational ease, we have omitted the dependence on $\mathbf{q}$. The Hessian matrix of $\hat{H}^0$ can be expressed as

$$\nabla^2 \hat{H}^0 = \begin{bmatrix} a_1 + a_2 + \dots + a_M & a_2 + \dots + a_M & \dots & a_M \\ a_2 + \dots + a_M & a_2 + \dots + a_M & \dots & a_M \\ \dots & \dots & \dots & \dots \\ a_M & a_M & \dots & a_M \end{bmatrix}.$$

For any $\mathbf{q} \in \mathcal{Q}$, we need to show $\nabla^2 \hat{H}^0 \succeq \epsilon I$ for some constant $\epsilon > 0$. As $\mathcal{Q}$ is compact, it suffices to show that the smallest eigenvalue of $\nabla^2 \hat{H}^0$ is strictly positive for any $\mathbf{q} \in \mathcal{Q}$. Notice that the diagonal matrix $\text{diag}\{a_1, a_2, \dots, a_M\}$ is *congruent* with $\nabla^2 \hat{H}^0$, as

$$\nabla^2 \hat{H}^0 = \begin{bmatrix} 1 & 1 & \dots & 1 \\ 0 & 1 & \dots & 1 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 \end{bmatrix} \begin{bmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & a_M \end{bmatrix} \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 1 & 1 & \dots & 1 \end{bmatrix}.$$

By Sylvester's Law of Inertia, the number of positive entries in $\text{diag}\{a_1, a_2, \dots, a_M\}$ is the same as the number of positive eigenvalues of $\nabla^2 H^0$ (Schultz 2011). Hence, for any $\mathbf{q} \in \mathcal{Q}$, the smallest eigenvalue of $\nabla^2 H^0$ is positive if every one of the a_i 's is positive, which leads to condition (ii) stated in the theorem.

EC.2. The Monotone Shortest Path Algorithm to Determine $\mathbf{q}^*$

We discuss the modified shortest-path algorithm used to determine $\mathbf{q}^*$ in the multi-demand model. Note that with multiple demands, it is possible to have $c_i < c_j$ for $i < j$, i.e., later arrivals may cost more than earlier ones. For example, if vessels i and j belong to the same shipping service and vessel j departs one week later, then vessels i and j have the same freight rate but c_j includes a higher inventory cost incurred at the origin. Consequently, $d_{ij} = w_{i,j}(y_{ij})$ in the graph $\mathcal{G}$ may *not* be positive. Therefore, when searching for the monotone shortest path, we would have to adopt the framework of Bellman-Ford Algorithm (see, for instance, Butler 2012), rather than the well-known and more efficient Dijkstra's Algorithm.

ALGORITHM 3. Initialize the graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$: Compute the parameters (y_{ij}, d_{ij}) for all $(i, j) \in \mathcal{E}$. For all $v \in \mathcal{V}$ and $i = 1, \dots, M + 1$, let $\text{dist}(v, i)$ represent the minimum distance from vertex v to vertex E that uses i or fewer edges. Set $\text{dist}(E, i) = 0$ for all i .

for all $v \neq E$ **do**

$\text{dist}(v, 1) := d_{v,E}$

$\text{succ}(v) := E$

end for

for $i = 2$ **do** $M + 1$

for all $v \neq E$ **do**

$\text{dist}(v, i) := \text{dist}(v, i - 1)$

for all $(v, x) \in \mathcal{E}$ **and** $x \neq E$ **do**

if $d_{v,x} + \text{dist}(x, i - 1) < \text{dist}(v, i)$ **and** $y_{v,x} < y_{x, \text{succ}(x)}$ **then**

$\text{dist}(v, i) := d_{v,x} + \text{dist}(x, i - 1)$

$\text{succ}(v) := x$

end if

end for

end for

end for

return $\text{dist}(O, M + 1), \text{succ}$

The above procedure is the standard Bellman-Ford algorithm except that we search only among the possible paths with y_{ij} strictly increasing. The output $dist(O, M+1)$ equals the minimum value of the objective function (13) and the output $succ$ keeps track of the vertices along the monotone shortest path, i.e., the optimal portfolio. The optimal shipping quantities can then be determined by the values of y_{ij} along the optimal path. The value of each y_{ij} can be easily computed from solving equation (20) or (21), whereas d_{ij} can be evaluated from our closed-form expressions. Given the parameters (y_{ij}, d_{ij}) , it is easy to see that the complexity of finding the monotone shortest path is $O(M^3)$. Therefore, the portfolio selection can be solved efficiently via the proposed dynamic programming-based method.

EC.3. The Impact of Lost Sales

To illustrate the impact of lost sales on the optimal portfolio, we consider four shipping services providing deterministic arrival times $\mathbf{T} = (5, 6, 7, 8)$ with freight rates $\mathbf{r} = (24, 20, 15, 11)$. Set $t_0 = 0$, $N = 14$, $h^D = h_i^B = .1$ and the price schedule as $p(t) = 61 - \frac{30}{7}(t - 1)$ for $1 \leq t \leq 15$. D follows a gamma distribution with $\mu_D = 100$ and $cv_D = 1.2$.




Figure EC.1 Effect of lost sales on the optimal portfolio

Figure EC.1a shows that with no lost sales, i.e., $1 - \beta = 0$, all four services are used, but more shipments go to the fastest service as the fraction of lost sales $1 - \beta$ increases. The slowest service (i.e., service 4) is the first to be eliminated from the portfolio when β drops to 0.4, and then services 3 and 2 are ruled out in turn. In the example, only the shipments via service 1 will arrive in the first period of the season and therefore be sold at the full price, whereas other services will lead to lower selling prices and possibly the loss of customers. In Figure EC.1b, we compare the

expected profit without diversifying with late arrivals (i.e., using only service 1) and that generated by the optimal portfolio. Because service 1 arrives in the first period, only using service 1 leads to an expected profit constant in β . When there are no lost sales (i.e., $1 - \beta = 0$), compared with supplying all shipments in the first period, diversifying with late arrivals can improve the profit by 20%. However, this improvement becomes insignificant when the fraction of lost sales is close to one. Therefore, an important caveat is that while the benefit of using slow services arises from time-sensitive prices, shippers should also pay attention to how unsatisfied demand is sensitive to time.