The impact of slow ocean steaming on delivery reliability and fuel consumption

Chung-Yee Lee [a,*], Hau L. Lee [b], Jiheng Zhang [a]

[a] Department of Industrial Engineering and Logistics Management, The Hong Kong University of Science and Technology, Hong Kong, China
[b] Graduate School of Business, Stanford University, Stanford, California 94305, United States

ARTICLE INFO

ABSTRACT

With a dominant volume of global transportation being conducted by sea, ocean container transport greatly impacts the global economy. Since sea vessels are drastically more fuel efficient when traveling at lower speeds, *slow steaming* has become a widely adopted practice to reduce bunker costs. However, this leads to a longer transportation time, which together with the unpredictability of the delay has been a big challenge. We propose a model to quantify the relationship among shipping time, bunker cost and delivery reliability. Our findings lead to a simple and implementable policy with a controlled cost and guaranteed delivery reliability.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

Globalization has extensively restructured the global supply chain in today's world economy. Products can be mass assembled at a site, with components shipped from multiple locations, and then distributed worldwide. This has been made possible, to a large degree, by the development of cost-effective and timely ocean shipping services which make the frequent exchange of on a large scale affordable. Ocean container transport plays an important role in global supply chains by connecting supply, manufacturing and distribution around the world.

Nevertheless, the ocean transportation industry is facing huge challenges. First, the fuel price has increased significantly over the years. Viewing the supply chain as a whole, the cost of bunker fuel is in one way or another shared by all those involved so that the direct monetary metric is important for everyone. Second, customers such as shippers and freight forwarders are increasingly demanding on-time delivery. Delivery reliability is of great importance for retailers who need to put products on their shelves, and for manufacturers who need to maintain the material flow to keep the factory running. According to Page (2010), in a new round of contract negotiations between shippers and a range of liners, priorities have been made clear and simple: service, reliability and price, in that order. In response, liners have started making firmer commitments to their customers to provide on-time delivery. Most contracts in this industry are set through party-to-party negotiation, and so the demand, price and reliability do not exhibit high elasticity as they do in airline revenue management. In this study, therefore, we focus on the two key performance metrics of bunker fuel cost and delivery reliability, without delving into the price/quality/demand issue.

We now describe the two challenges – delivery reliability and fuel cost – in more detail. Delivering on time is not easy even with today's advanced nautical technology. Late deliveries can result from port congestion, inefficient port operations, extreme weather conditions, machine breakdowns and other factors. Among them, port congestion has become the most

---

* Corresponding author.

important contributing factor. According to Notteboom (2006), who used survey data on the East Asia–Europe route, unexpected waiting before berthing as a result of port congestion contributed to 65.5% of the schedule unreliability. In general, Notteboom (2006) singled out randomness in the waiting and service times at ports as the major obstacle to making deliveries on time. Our own analysis of data from Orient Overseas Container Line (OOCL) supports his findings.

The variance is due primarily to the number of box moves, but port congestion and varying productivity at terminals also contribute. The significance of the variance can be observed from the histogram (depicted in Fig. 1) of port times at the Hong Kong port for all vessels on OOCL's trans-Pacific service route during 2009–2011. Port congestion is a result of the tremendous growth in the demand for container transport which has outstripped the growth of container handling capacity at ports. The industry expects that port congestion will remain an exogenous factor. In addition to the waiting time caused by congestion, the vessel processing time which depends on the number of boxes to be offloaded and loaded also has a large variability. The resulting variability of total port time has become a huge challenge for liners hoping to deliver boxes on schedule.

In the case of fuel cost, the fuel consumption of sea vessels depends heavily on the steaming speed. Analysis of the OOCL data shows that increasing the speed by a couple of knots burns almost 50% more fuel per unit of distance traveled (see Fig. 2). Ocean vessels could make up for delays by properly speeding up, but this action generates much higher fuel costs. According to the World Shipping Council WSC (2008), fuel costs represent as much as 50–60% of total operating costs. In the past when the fuel price was much lower, the steaming speed was not a big issue. However, as the fuel price continues to climb, the whole ocean transport industry has been forced to move from 23–25 knots steaming to 20–22 knots steaming, or even extra slow steaming at 17–19 knots. Slow steaming lengthens the round-trip time by 10–20% depending on the service route and port times along the route. Though this is undesirable, shippers usually accept it following price and service
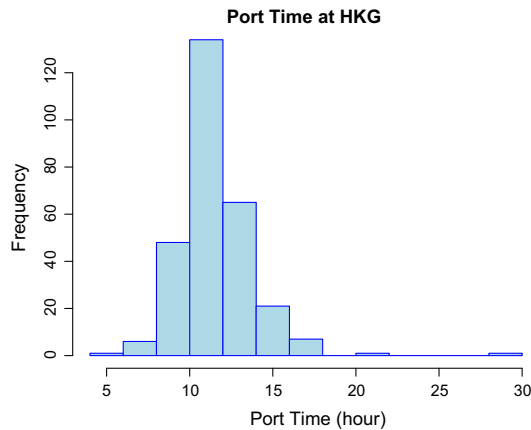

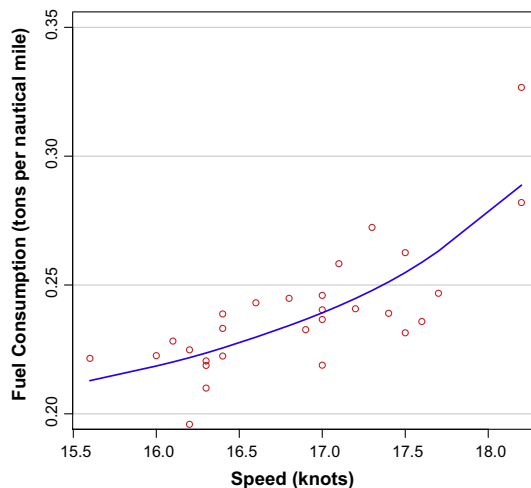**Fig. 1.** Histogram of total port times at the port of Hong Kong.

Fig. 2. Fuel efficiency of sea vessels (data from OOCL 8000-TEU vessels sailing from LGB to KHH).

negotiations. According to OOCL, 90% of its Euro-Asia routes and 60% of its Trans-Pacific routes have adopted slow steaming. A lower base speed now yields more room to speed vessels up. In fact, the extra fuel cost of speeding up by 1 knot from 16 knots is much less than that of speeding up by 1 from 22 knots, since the fuel consumption function is sharply convex. Freight rates, however, have stayed roughly the same as the increased fuel price cancels out the savings on fuel. Furthermore, customers do not see any improvement in delivery reliability as liners remain reluctant to speed up even when the base speed is lower than before. A central question is how can ocean carriers make use of slow steaming for fuel savings while at the same time improving delivery reliability?

It has been an industry standard for liners to call each port on a service route at a fixed frequency (normally once a week). Thus slow steaming may require liners to deploy more vessels on a given route. OOCL vessels used to sail at an average speed of 22.4 knots, leading to a round-trip time of 56 days on a particular service route. To maintain weekly service to ports along that route, they deployed 8 vessels. Since adopting slow steaming, their vessels have been sailing at an average speed of 20.2 knots and the round-trip time has been extended to 63 days. Consequently, they have to deploy one more vessel to maintain the same service frequency. Even with the additional vessel, slow steaming still helps substantially reduce the total operating cost due to the huge savings on fuel. Using estimates from AlphaLiner, Bonney (2010) has shown that with the bunker fuel price at $500 per ton, slow steaming reduces the total operating costs on long-haul loops by 5–7%. The annual savings could reach $15 million to $20 million for a typical Asia–Europe route using 8500-TEU ships. Bunker fuel costed between $1000 and $1500 per ton[1] in the year of 2014. Slow steaming has drawn widespread attention as it is both a cost saver and environmentally friendly (see Cariou, 2010 for a quantitative estimation of the reduction in emissions due to slow steaming).

There are three layers of decision making in the liner container business: strategic, tactical and operational. At the strategic level, liners deal with long-term planning issues such as vessel procurement. At the tactical level, they deal with issues such as vessel routing and schedule planning. Once routing and scheduling are sorted out, at the operational level, they are interested in finding ways to manage their operations better, for example, by adopting slow steaming on a voyage. Our paper deals with the operational level decision. In particular, we study delivery reliability and fuel consumption by incorporating port time randomness into a stylized yet practical model. We do not concern ourselves with the decision analysis on purchasing new vessels, which is a long-term investment, and chartering, which serves as an option when a liner is incapable of or decides not to make a long-term investment. We also ignore the demand and its dependence on price, service quality and competition in our analysis, since demand does not depend strongly on these factors as mentioned above.

In general, operations research can also help improve liner shipping services by offering insights into two key aspects. The first aspect is the design of the service routes and delivery schedules. A service provider decides a route consisting of a sequence of ports and a fleet consisting of several vessels of a certain size to operate on this route. The selection of routes and fleet depends on the service frequency, which in turn depends on the demand at the ports along the route. Once the route and fleet size are fixed, the liner needs to set the delivery schedule, i.e., the dates on which its vessels are scheduled to arrive at the ports along the route. The delivery schedules are normally published quarterly. Designing the delivery schedule requires quantitatively studying the relationship among the shipping time, delay and fuel consumption. This is one of the focuses of our paper. The second key aspect is the operational decision of when to speed up a vessel upon the realization of random port times (and potential delays) on service routes. Typically, this decision also concerns the speed at which a vessel should sail toward downstream ports. If the observed delay is significant, it might not be possible to make up for all the time lost. In this paper, we will also incorporate the operational decision of when to switch from a low speed to a high speed.

Our paper contributes to the literature by providing a better understanding of the impact of random port operating times (including waiting times at ports) on delivery reliability and fuel consumption cost. We propose a model and use it to analyze the inter-relationships of delivery reliability, fuel cost, and transportation lead times by giving some bound estimations. The significance of our results is to provide some implementable policy to achieve certain service quality (measured by delay probability) while keeping the fuel consumption under control.

After reviewing the literature in Section 2, we propose a model in Section 3 to capture steaming speed as the decision variable and port time as the source of uncertainty. Section 4 studies the impacts of slow steaming on the delay probability, variance of delivery time, fuel consumption and shipping time. Section 5 discusses how these impacts affect both the liners and the shippers. A simulation study based on data collected from OOCL is also provided. We conclude our paper in Section 6.

2. Literature review

Operations research in ocean transportation has traditionally focused on terminal operations. Steenken et al. (2004) provided a comprehensive survey on the state of the art of operations at a container terminal. Stahlbock and Voß (2008) reviewed more recent works on (a) the design of terminal structure, handling equipment, human resources and supporting systems, and (b) operation and optimization of terminal logistics.

Another line of research relates to vessel routing to minimize transportation cost, which is a relevant problem due to rising fuel prices. Fagerholt et al. (2010) studied an optimization model based on a shortest path problem on a directed acyclic

---

[1] http://www.bunkerworld.com/prices/.

graph for minimizing fuel consumption subject to the constraint that deliveries at each port on a predetermined service route must be made within certain time windows. Notteboom and Vernimmen (2008) studied how shipping liners have adapted their liner service schedules (in terms of commercial speed, the number of vessels deployed per loop, etc.) to deal with increasing bunker costs. A simple cost model was used to simulate the impact of bunker cost changes on operating costs in liner services. Recently, Qi and Song (2012) proposed simulation-based stochastic approximation methods to design an optimal delivery schedule for a given route to minimize the total expected fuel consumption while considering the uncertainty of port times. They pointed out that the relationship between fuel consumption and liner service design has not drawn enough academic attention, even though it has been a major concern of the shipping liners. Instead of using simulation, we use an analytical model to estimate the inter-relationship of delivery reliability, fuel cost and transportation lead time.

For the design of service routes, Rana and Vickson (1991) formulated a mathematical programming model to determine the optimal sequence of ports of call, service frequency, and the number of cargo units to be transported between each pair of ports by each ship. Fagerholt (1999) studied the problem of determining an optimal fleet (the type of ships and the number of each type) in a special network where all cargos are transported from a set of production ports to a single depot. Ting and Tzeng (2003) developed a dynamic program for scheduling decisions like cruising speed, quay crane dispatching, and rough schedule arrangements. However, they did not consider port congestion. Agarwal and Ergun (2008) proposed a model for the design of efficient service routes given a set of demands to be transported and a set of ports to be served. They presented a mixed-integer linear program to solve the ship scheduling and cargo routing problems simultaneously. Leveraging this model and game theory concepts, Agarwal and Ergun (2010) further designed a mechanism to guide the carriers in an alliance to pursue an optimal collaborative strategy. The mechanism provides side payments to the carriers, as an added incentive, to motivate them to act in the best interest of the alliance while maximizing their own profits. Brouer et al. (2014) proposed a base integer programming model and benchmark suite for liner-shipping network design. Álvarez (2009) presented a model and an algorithm to jointly determine the optimal routing and deployment of a fleet of container vessels. Readers interested in the literature on ship scheduling and routing may turn to the reviews by Ronen (1983, 1993) and Christiansen et al. (2004). Most works in this line of research either assumed deterministic port times, or did not consider the possibility of speeding up. Wang and Meng (2012a) studied the problem of designing a tactical-level liner schedule. They developed a mixed-integer non-linear stochastic programming model for the problem by minimizing the expected bunker cost and other related cost while maintaining a required service level in terms of transit times. Christiansen et al. (2013) and Meng et al. (2014) provided a review and outlook for ship routing and scheduling. In our paper, we are given a service route and our purpose is to analyze the relationship of reliability and transportation lead time.

Recently, in an expositional paper, Fransoo and Lee (2013) identified several key questions and issues relevant in supply chain management: the coordination of container shipments across the container supply chain, pricing and risk management in the container supply chain, competition between ports, carriers and container terminals, and capacity management in the container supply chain. Furthermore, they explicitly mentioned the need to study the extent to which operational variability exists and how it affects decision making. Indeed, how to manage time is an important issue in contemporary liner services, since significant waiting times and delays at ports put pressure on schedule reliability.

A survey by Vernimmen et al. (2007) revealed that over 40% of the vessels deployed by liners worldwide arrive one or more days behind schedule, despite liners' claims that most of their container ships operate on fixed weekly schedules. A recent survey by Drewry (2010) gave more detailed statistics by carrier, trade and service from December 2005 to June 2010. The industry average of delay was found to range from 32% to 54%. The delays not only caused complaints, but also imposed real costs on the shippers and their customers. Vernimmen et al. (2007) presented a case study to illustrate the impact of schedule unreliability on the level of safety stock needed by a manufacturer who sources spare parts from overseas. Their analysis showed that an improvement in schedule reliability can lead to significant cost savings for the manufacturer. Notteboom (2006) identified the causes of schedule unreliability, and discussed a wide array of measures and planning tools that liners deployed to maximize schedule reliability. However, the paper did not provide any quantitative result on the tradeoffs between these measures and schedule reliability.

Slow steaming has began attracting research attention in recent years. Notteboom and Cariou (2011) provided data from mid 2008 to late 2010 to show that (i) slow steaming has become an industry-wide trend; (ii) slow steaming leads to longer shipping times but much lower fuel cost; and (iii) slow steaming does not lead to sharing of savings on fuel cost and the fuel surcharge stays more or less the same as that before slow steaming. The second point further supports the quantitative insight we have obtained using our modeling approach. Ronen (2011) developed a cost model to study the optimal speed and the number of vessels needed for maintaining a service frequency while minimizing the total cost, including bunker cost, vessel fixed cost, and other operating cost. Wang and Meng (2012b) used historical operating data of a global shipping liner to study the relationship between bunker consumption and sailing speed. They formulated a mixed-integer nonlinear programming model to investigate the optimal speed. They also provided an efficient approximation method to obtain a nearly optimal solution. Some studies examined the impact of slow steaming from the supply chain's viewpoint, thus the cost model included logistics-related costs such as transit inventory holding cost. For example, Psaraftis and Kontovas (2010), and Cariou (2011) discussed the tradeoff between the costs (including logistics-related cost, vessel operating cost) and potential benefits (including bunker cost and CO_2 emission reduction) of slow streaming. Psaraftis and Kontovas (2013) provided a comprehensive survey on models involving steaming speed as a key decision variable. In those models, the fuel consumption (which is proportional to CO_2 emission) and revenue are the main concerns. However, the impact of port congestion and the randomness of port time, an increasingly important factor, on fuel consumption has not been analyzed in those models.

Corbett et al. (2009) used a speed optimization in the profit function to study the impact of fuel tax on CO_2 reduction. Note that fuel tax will affect the profit function which will then indirectly affect the optimal speed and CO_2 reduction. Cariou and Cheaitou (2012) argued that the speed limit that the European Union (EU) is thinking of imposing on all ships entering EU ports may ultimately generate more emission, and may be suboptimal. In the above literature, they either focus on fuel consumption (which is proportional to CO_2 emission) or the trade-off between bunker consumption and sailing speed, our paper mainly aims at providing quantitative estimates of delivery reliability, fuel consumption and transportation lead time and hope to provide managerial insights and useful tools for the industry.

We want to briefly mention a stream of literature on train and metro systems. Feng et al. (2011) studied the maximum operation speeds of metro trains for both energy saving and transport efficiency improvement. Li and Lo (2014a) developed an optimization method to improve the operations of metro rail system. Li and Lo (2014b) proposed a dynamic train scheduling and control framework for metro rail system by solving a convex optimization model to improve its energy saving performance. There are some difference between the metro systems and ocean transportation. Due to the short distance between stops in subway systems, synchronizing accelerating and braking trains is important. The time trains stop at each station exhibit relatively less variability, which are not considered in the studies we found.

While the time dimension of container shipping is critical to the service level, there have been limited quantitative studies in the literature. The adoption of slow steaming and the emphasis on delivery reliability point to the importance of studying the relationships among steaming speed, shipping time, fuel costs and delivery reliability.

3. Problem statement and notations

We introduce a model that incorporates the randomness of port times and uses steaming speed as the decision variable to study the relationship among the leading performance metrics of fuel consumption, delivery reliability (measured by the variance of the actual delivery time), and probability of delay.

3.1. Port time

Consider a liner that deploys a fleet on a fixed service route, which consists of a sequence of ports that each vessel in the fleet needs to visit. We number the ports by $0, 1, \ldots, K$. For example, if a route consists of *Rotterdam, Singapore, Hong Kong, Shanghai, Hong Kong, Rotterdam*, then K is 5 with $0 = Rotterdam, 1 = Singapore, 2 = HongKong, 3 = Shanghai, 4 = HongKong, 5 = Rotterdam$. A vessel arriving at a port may experience some waiting time based on the congestion level of the port. It will then be assigned to a terminal at which it will load and unload container boxes. According to OOCL's historical data, the total amount of time (both waiting and service) that a vessel spends at a port was quite random. Fig. 1 depicts the histogram of the port times experienced by all vessels deployed on a trans-Pacific service route of OOCL at the port of Hong Kong in the past three years. Let W_k be the *port time* experienced by a vessel at port k, including both the *waiting time* and the *service time*. In other words, W_k is the amount of time from the arrival at port k till the departure from port k.

3.2. Steaming speed

Let D_k denote the distance from port $k - 1$ to k. For planning purposes, liners need to set a default steaming speed v in order to make a timetable of delivery at all ports along the service route. In real time, say when a vessel is about to leave port k, there is also a decision to be made about what speed should be applied in navigating to port $k + 1$, based on how much delay has been accumulated so far. By speeding up, a vessel can normally recover part of the time lost, at the expense of much higher fuel cost. For example, OOCL has found that for the trip from Los Angeles (LGB) to Kaohsiung (KHH), the total consumption of heavy fuel by the main engine is 2017 tons when the vessel is traveling at an average speed of 18.2 knots, and 1493 tons when the vessel is traveling at an average speed of 17.0 knots. A 1.2-knot increase in sailing speed increases fuel consumption by more than 30%. Shipping liners often do not have any model to estimate the tradeoff between the extra cost of speeding up and the benefits of making a delivery on time. In many situations, they simply choose not to speed up and let the schedule slip. For this purpose, we try to give a bound estimate of the extra cost of speeding up by comparing a flexible slow steaming strategy and a fast steaming strategy and quantifying the effects of these strategies for fuel consumption, reliability and transit time, when accounting for unpredictability in delays.

3.3. Delivery Schedule

Liners normally work out a delivery schedule once a service route has been chosen and publish it to their customers. The schedule basically tells their customers or partners when container boxes will be delivered at each port. In practice, the liners set a scheduled time $s_k(v)$ between ports (time between departing port $k - 1$ and departing port k) by applying a calculation that considers the distance, the estimated waiting time, the amount of loading and unloading (mainly for the purpose of estimating the expected service time), and a buffer time. The current practice is to estimate the number of "moves" (a "move" is defined as the action of loading or unloading one container) and the speed of the cranes (measured by the number of moves

per hour) at the terminal to arrive at an "estimated service time". A certain amount of time called "port contingency" is then added as a buffer to deal with randomness. Essentially, liners are trying to set the schedule using the following formula

$$s_k(v) = \frac{D_k}{v} + \mu_k + \beta_k \sigma_k, \tag{1}$$

where v is the planned speed as discussed before, μ_k and σ_k are the mean and standard deviation of port time W_k, and the coefficient β_k is a parameter chosen to denote the "port contingency", which plays a role as a buffer to mitigate any delay. This is in the same spirit as many inventory policies, but we shall keep in mind that the schedule overestimates the actual shipping time. Instead of estimating important quantities in an ad hoc way, we can take advantage of historical data and use a quantitative model to achieve better management goal. Setting the time of leaving from port 0 to 0, the scheduled time to *depart from* port k is

$$S_k(v) = \sum_{i=1}^{k} s_i(v), \quad k = 1, 2, \ldots, K. \tag{2}$$

Tracking when a vessel departs from a port gives a better estimate of when shippers can expect their deliveries than using the arrival time at the port. The dependency of the delivery timetable on the chosen default steaming speed v is clear.

It is an industry practice to make the round-trip time multiple of weeks so that each port will be called weekly. An arbitrarily chosen speed may not satisfy this constraint. We first decide the number of weeks for the total round-trip time, and then use the formulae to determine the speed. If the resulting speed falls out of the range of the vessel's physical speed, then we need to consider modifying the round-trip time to be some other multiple of weeks so that feasible speeds can be obtained.

3.4. Delay

We start by assuming that the vessels can only navigate at either low speed v_L or high speed v_H. We will later discuss the general case where the speed can be continuously adjusted. Due to the randomness in W_k, a vessel may not always meet the schedule. A Markovian model can describe the evolution of delay as a vessel navigates from port to port on a service route. We will apply this model to compare the amount of delay at each port under two possible strategies.

The first strategy, called *fast steaming*, is for a vessel to sail at the high speed v_H throughout its entire journey. This strategy was the practice before slow steaming became popular, and serves as a benchmark for analyzing the impact of slow steaming. Let Y_k denote the amount of delay upon leaving port k. Each vessel on the service route starts from the home port (port 0). After visiting all the ports on the route, it returns home. For simplicity, we assume that $Y_0 = 0$. See Remark 1 for discussion on the case where $Y_0 > 0$. The delay at downstream ports evolves according to the following Markov chain:

$$Y_k = Y_{k-1} + \frac{D_k}{v_H} + W_k - s_k(v_H), \quad k = 1, 2, \ldots, K. \tag{3}$$

Using (1), we have

$$Y_k = Y_{k-1} + W_k - \mu_k - \beta_k \sigma_k, \quad k = 1, 2, \ldots, K. \tag{4}$$

Due to the randomness, Y_k may be negative meaning that the vessel may get ahead of the schedule.

The second strategy, called *flexible slow steaming*, is for the liner to plan its schedule based on the low speed v_L and for the vessels to sail at this speed by default. However, if there is a delay, the vessel is allowed to sail at the high speed to make up for time lost. To simplify the notation, let $c = \frac{v_H - v_L}{v_H v_L}$. It is clear that cD_k is the maximum amount of time lost a vessel can make up for by switching to the high speed on the trip from port $k - 1$ to port k. When the delay upon leaving port $k - 1$ is less than cD_k, the vessel will only apply the high speed over part of the distance D_k to make up for the exact amount of time lost. Let X_k denote the amount of delay upon leaving port k in this case. Similar to the previous case, we have $X_0 = 0$ and the following induction

$$X_k = (X_{k-1} - cD_k)^+ - X_{k-1}^- + \frac{D_k}{v_L} + W_k - s_k(v_L), \quad k = 1, 2, \ldots, K, \tag{5}$$

where $a^+ = \max(a, 0)$ and $a^- = \max(-a, 0)$ for any $a \in \mathbb{R}$. Again, applying (1), we have

$$X_k = (X_{k-1} - cD_k)^+ - X_{k-1}^- + W_k - \mu_k - \beta_k \sigma_k, \quad k = 1, 2, \ldots, K. \tag{6}$$

Remark 1. Our comparison results still hold if $X_0 = Y_0 = a$ for some constant a. However, the ports are visited in a cyclic way as vessels embark on consecutive journeys. It could happen that cumulative delay during a few journeys becomes to prominent. In this case, disruption recovery measures will be taken. Li et al. (in press) provided a study on disruptions recovery.

By writing the Markov chains (4) and (6), we allow a vessel to leave a port earlier than scheduled if it is ahead of the schedule. This is in fact a standard industry practice now for two reasons. First, port delay or earliness is measured in hours, however container boxes must go through several processes including a security check, and normally sit in the terminal yard for days before the scheduled pickup time. Second, terminals are extremely busy and cannot afford to let vessels occupy precious quay space just to wait for a few boxes. So when the service is completed, the vessel must set sail, sometimes (although rarely) leaving a few scheduled boxes behind. We should acknowledge that vessels may further slow down when getting ahead of schedule by too much. However, we are not able to obtain any analytical result when exploring this dimension of flexibility. In this model, liners will not speed up to get ahead of the schedule in order to hedge against any potential delays at downstream ports. In other words, liners are myopic. This makes sense, since fuel consumption constitutes the major cost and vessels become much less fuel-efficient at high speeds. It is also not necessary to get ahead if the planned schedule has already factored in potential delays with buffers. In fact, due to the lack of a quantitative model to estimate fuel consumption and fluctuating fuel prices, some liners that have adopted slow steaming avoid speeding up even when their ships are delayed. To manage the risk of unexpectedly long waiting times, they find it more cost effective to use the coefficients β_k's in the scheduling formula than to speed up. In ocean transportation, being on time is more critical than reducing the lengthy shipping time. Building on this model, we will discuss in the next section how liners' service quality is affected by their choice of strategy and the parameters β_k's.

4. The impact of slow steaming

In this section, we discuss the impact of slow steaming on two measures of service quality in Sections 4.1 and 4.2. We will then analyze the fuel consumption in Section 4.3. There are various measures of service quality. One of them is the delay (or on-time) probability. Note that an alternative measure could be the expected delay, but the industry currently uses the delay probability instead of its expectation due to the former's ease of implementation (see Platt et al., 1997 for a similar discussion of constrained service levels in inventory management). The probability that the delivery is delayed by more than a tolerable threshold is therefore of interest. Another important measure of service quality is the accuracy of the delivery time, since shippers want their deliveries to be made as close as possible to the promised date. This offers "predictability" for shippers in managing their businesses. Predictability hinges on the variance of delivery time.

4.1. Delay probability

Consider the service quality measured by the probability that the delay exceeds a certain amount $s \geqslant 0$. We want to compare the service quality under the flexible slow steaming strategy, i.e., $\mathbb{P}(X_k > s)$, with that under the constant fast steaming strategy, i.e., $\mathbb{P}(Y_k > s)$ when the schedule is set using (1) with the same β_k's. Note that when $s = 0$ the probabilities are simply the delay probabilities.

Proposition 1. *Assume that port times are independent of the delays. For any constant s, the inequality* $\mathbb{P}(X_k > s) \leqslant \mathbb{P}(Y_k > s), k = 1, \ldots, K$, *holds for the Markov chains* $\{X_k\}_{k=1}^K$ *and* $\{Y_k\}_{k=1}^K$ *defined in (4) and (6).*

The proof can be found in Appendix A. The intuition behind the result is that the Markov chain $\{X_k\}$ is always stochastically smaller than $\{Y_k\}$. Note that this result does not need any assumption on the shape of the distributions of port times W_k, and is actually independent on the choice of β_k in (1). In fact, (1) can be replaced by a more general rule,

$$s_k(v) = \frac{D_k}{v} + A_k(W_k), \tag{7}$$

as long as A_k is a deterministic function of W_k. In other words, if the delivery schedule is set using formula (7), slow steaming with the flexibility to speed up will always yield a better service quality as measured by the probability that the delay exceeds a certain amount than constant fast steaming.

4.1.1. Service quality under the fast steaming policy
The tractability of the Markov chain $\{Y_k\}$ suggests an iterative way of setting the β_k's so that the delay probability can be reduced to less than $1 - q_k$ for all ports $k = 1, \ldots, K$. By Proposition 1, $1 - q_k$ serves as a lower bound if we apply the flexible slow steaming strategy.

The evolution of the Markov chain $\{Y_k\}$ yields

$$Y_k = \sum_{i=1}^k (W_i - \mu_i) - \sum_{i=1}^k \beta_i \sigma_i. \tag{8}$$

For technical reasons, assume that the port times are independent and normally distributed. Although the histogram (c.f. Fig. 1) of port times exhibits a unimodal shape that resembles the normal distribution, the port times are not exactly normally distributed as statistical tests such as the Lilliefors test would show. One possible reason is that the sample size of a few hundred is too small. On trans-Pacific service routes, the round-trip time is about a month and a half. Over a period of three years, only a few hundred data points can be accumulated. Even though data exist describing a much longer period,

data from too far back in the past may not represent the current situation at ports. Approximating the random variables using the normal distribution is the simplest way to incorporate both the mean and variance of port times into the analysis. In fact, it will be shown that the delay probability obtained based on the assumption of a normal distribution is quite robust. Based on this assumption, we further have

$$Y_k \overset{d}{=} \mathcal{N}\left(0, \sum_{i=1}^{k} \sigma_i^2\right) - \sum_{i=1}^{k} \beta_i \sigma_i, \tag{9}$$

where $\mathcal{N}$ represents a normally distributed random variable with the first parameter denoting the mean and the second one denoting the variance. This provides an inductive way of setting the parameters β_k's for $k = 1, \ldots, K$. Suppose that the delay probability at port k must be less than $1 - q_k$, which is chosen by liners based on their evaluation of the importance of service quality at port k. By (9), $\mathbb{P}(Y_1 > 0) = 1 - q_1$ basically requires that $\mathbb{P}(\mathcal{N}(0, 1) \leqslant \beta_1) = q_1$. Let z_p denote the $100p$th percentile of a standard normal random variable for any $p \in [0, 1]$. Thus $\beta_1 = z_{q_1}$. This procedure can be continued inductively. Suppose $\beta_1, \ldots, \beta_{k-1}$ have been determined based on $q_1, \ldots, q_{k-1}$. Based on (9), the general formula that connects the parameters is

$$\sum_{i=1}^{k} \beta_i \sigma_i = z_{q_k} \left(\sum_{i=1}^{k} \sigma_i^2\right)^{1/2}. \tag{10}$$

The following formula can then be used to determine β_k based on q_k:

$$\beta_k = \frac{z_{q_k} \left(\sum_{i=1}^{k} \sigma_i^2\right)^{1/2} - \sum_{i=1}^{k-1} \beta_i \sigma_i}{\sigma_k}. \tag{11}$$

This formula describes a procedure to set the parameter β_k inductively, but alternatively, (10) suggests that

$$\beta_k = \frac{z_{q_k} \left(\sum_{i=1}^{k} \sigma_i^2\right)^{1/2} - z_{q_{k-1}} \left(\sum_{i=1}^{k-1} \sigma_i^2\right)^{1/2}}{\sigma_k}. \tag{12}$$

This formula prescribes a way of managing the service quality q_k by setting the parameter β_k in schedule planning. It also has a few implications. First, note that the choice of β_k depends on the following factors:

(i) the standard deviation σ_k of port time at port k,
(ii) the accumulated variance of all port times up to port $k - 1$ and that up to port k, and
(iii) the desired levels of service quality q_{k-1} and q_k at ports $k - 1$ and k.

If the service quality at the previous port is quite high, i.e., $z_{q_{k-1}}$ is large, then β_k can be relatively small. It is worth pointing out that β_k does not necessarily need to be bigger than 0, meaning that a positive buffer or "port contingency" is not always needed. We can easily imagine a case where $z_{q_k} \left(\sum_{i=1}^{k} \sigma_i^2\right)^{1/2}$ is less than $z_{q_{k-1}} \left(\sum_{i=1}^{k-1} \sigma_i^2\right)^{1/2}$, yielding a negative β_k. This is somewhat counter-intuitive, but an important finding nonetheless. Reducing the extra buffer time can either reduce the planned delivery time, which will help make a liner's service more attractive, or allow more time at sea, which means a slower speed and savings on bunker costs. We will explain the bunker cost part in more detail in Section 4.3.

We conclude this section with a special case to illustrate the application of formula (12). Consider the case where all port times are i.i.d. with mean μ and variance σ^2. Suppose the same service quality is enforced at all ports, i.e., the delay probability must be less than $1 - q$. In this case, formula (12) yields $\beta_k = z_q \left(\sqrt{k} - \sqrt{k-1}\right)$. The extra buffer time, quantified by β_k, then decreases from port to port throughout the voyage. This may seem to contradict the general belief that it becomes more difficult to control the timing or to stay on schedule as a voyage progresses since the variance at a downstream port is the combined variances from all previous ports. But the truth as revealed in (12) is that the magnitude of β_k is determined by the difference between the accumulated variances up to port k and the accumulated variance up to port $k - 1$, and not by the total accumulated variance. Intuitively, the main reason that β_k decreases is that early arrival is allowed.

4.2. Variance of the delivery time

We now discuss the variance of delivery time under the two different strategies. As mentioned above, what shippers care most about is the "predictability" of the deliveries. This has also been emphasized by the carriers. According to Eivind Kolding, the former chief executive of Maersk Line, the leading company in the shipping industry, "Reliability is the new rate war; we need an end-to-end view on reliability." (c.f. Kolding, 2011).

Let T_k^H denote the delivery time at port k in the case where the high speed is applied all the time, and T_k^L denote the delivery time at port k in the case where the low speed is applied with the flexibility to speed up whenever there is a delay. As discussed in Section 3, the departure time is a better approximation of the delivery time, thus we use the departure time as the delivery time in the discussion in this section. It is clear that

$$T_k^H = S_k(v_H) + Y_k, \quad \text{and} \quad T_k^L = S_k(v_L) + X_k. \tag{13}$$

Note that T_k^H and T_k^L measure the time from leaving port 0 to when products are delivered at port k, but the time between any two ports k and k' can be easily computed. For simplicity, we will focus on the discussion of T_k^H and T_k^L. The question is whether the flexible slow steaming strategy can lead to a more punctual delivery time at each port. We show that the variance of the delivery time under the flexible slow steaming strategy is always smaller than that under the fast steaming strategy, meaning more accurate deliveries, though the expected shipping time is longer.

Proposition 2. *Assume that port times are independent of the delays. For the delivery times T_k^H and T_k^L as defined in* (13) *we have $Var(T_k^L) \leqslant Var(T_k^H)$, for all $k = 1, \ldots, K$.*

The proof can be found in Appendix A. It is worth pointing out that Proposition 2 does not make any assumption about the distribution of port times W_k either. Moreover, there is no need to use the same formula to set the schedule. This is different from the result in Proposition 1, where the same formula is required for creating the schedule. For example, instead of using the same formula (1) to set the schedule for both high and low speeds, we can use $\{\beta_k^H\}$ to set the buffer times for fast steaming and $\{\beta_k^L\}$ to set those for flexible slow steaming. In summary, the flexible slow steaming strategy is a robust strategy that reduces delivery variability.

4.3. Fuel consumption

Putting the benefits of improved service quality aside, an important reason for using slow steaming is the significant reduction in fuel consumption. We examined the fuel consumption data for the 8000-TEU vessels currently deployed on the trans-Pacific service route offered by OOCL. All vessels have the same cost characteristics. Let $m(v)$ denote the amount of fuel consumed per nautical mile for these sea vessels at the steaming speed v. Fig. 2 plots the fuel efficiency, measured by tons of heavy fuel consumed per nautical mile, versus the steaming speed for these 8000-TEU OOCL vessels deployed on their trans-Pacific service routes. We chose the longest leg on the route, between Long Beach (LGB) and Kaohsiung (KHH), since the estimates for average steaming speed and fuel consumption are more accurate for longer distances. The data points represent all LGB–KHH trips made between January and October of 2011, amounting to a sample size of about 30. The curve is the smooth fitting of the data using spline. The fitted curve shows a monotonically increasing trend. In fact, the increase is quite sharp, since the horizontal axis only ranges from 15.5 to 18.5 knots. Similar plots have been observed in other studies in the literature (Fransoo and Lee, 2013; Notteboom, 2006; Vernimmen et al., 2007; Rodrigue et al., 2009).

Denote by $D = \sum_{i=1}^{K} D_i$ the total travel distance on the service route. The fuel consumption per trip per vessel under the fast steaming strategy is simply $Dm(v_H)$. Under the flexible slow steaming strategy, however, deriving the fuel consumption per trip is more complicated, since not every nautical mile is travelled at the low speed under this strategy. Essentially, we need to estimate the distance that is navigated at the high speed. The distance from port $k-1$ to port k is D_k, and the distance navigated at the high speed depends on X_k^+ and can be written as $\frac{1}{c} \min(X_k^+, cD_k)$. It is not possible to perform an exact analysis on the above term, but an upper bound for its expectation can be estimated by comparing the Markov chains $\{X_k\}$ and $\{Y_k\}$.

Let G_{X_k} and G_{Y_k} denote the cumulative distribution functions of X_k and Y_k, respectively. The following upper bound estimation follows from Proposition 1.

$$\mathbb{E}[\min(X_k^+, cD_k)] = \int_0^{cD_k} [1 - G_{X_k}](x)dx \leqslant \int_0^{cD_k} [1 - G_{Y_k}](x)dx \leqslant cD_k \mathbb{P}(Y_k > 0).$$

If the parameters β_k are chosen to meet service level q_k as in the previous section, then

$$\mathbb{E}\left[\frac{1}{c} \min(X_k^+, cD_k)\right] \leqslant D_k(1 - q_k). \tag{14}$$

Since the analysis applies for all k, we have the following result.

Proposition 3. *Assume that port times are independent of the delays. The expected fuel consumption for the flexible slow steaming strategy is bounded from above by $m(v_L)\sum_{i=1}^{K} D_k q_k + m(v_H)\sum_{i=1}^{K} D_k(1 - q_k)$.*

This implies that the saving in fuel consumption per trip per vessel is at least $[m(v_H) - m(v_L)]\sum_{i=1}^{K} D_k q_k$ if the flexible slow steaming strategy is adopted. If the delay probability is less than 10% for all ports, i.e., $q_k = 0.9$ for all k, then the amount of fuel saved per trip would be $0.9[m(v_H) - m(v_L)]D$. Since the function $m(\cdot)$ increases with a fairly steep slope, the saving could be quite significant indeed.

Such an upper bound is quite useful in practice. In today's liner business, fuel has become increasingly large part of operating costs. Some liners, having adopted slow steaming, are reluctant to speed up even when they are falling behind the schedule. The main reason is the lack of quantitative insight on how much more fuel need to be burned if they speed up whenever they fall behind schedule. Our formula shows that the cost of speeding up, following the flexible slow steaming policy, is bounded by

$$[m(v_H) - m(v_L)]\sum_{i=1}^{K} D_k(1 - q_k), \tag{15}$$

which is lower than the savings from adopting slow steaming. Thus it can be affordable for the liners to speed up.

4.3.1. A general case with continuously adjustable speed

So far, the steaming speed has been limited to two choices, $\{v_L, v_H\}$. In this section, we show that if the steaming speed is continuously adjustable, i.e., the speed can be chosen from anywhere in the interval $[v_L, v_H]$, then the savings on fuel consumption can be even greater given that the efficiency function $m(v)$ is convex.

Note that it is always better to keep the speed constant rather than to vary it if the function $m(v)$ is convex. The convexity is a valid assumption based on our data analysis (c.f. Fig. 2 (RHS)). Suppose a vessel sails at speeds v_1 and v_2 (with $v_1 < v_2$) for a distance D, with v_1 being applied for a fraction $p \in [0, 1]$ of the distance D, and v_2 being applied for the remaining distance. Then fuel consumption for the distance D would be $pDm(v_1) + (1 - p)Dm(v_2)$. However, the same navigation time could be achieved at the constant speed

$$v = \frac{v_1 v_2}{(1-p)v_1 + p v_2} = \frac{p v_2}{(1-p)v_1 + p v_2} v_1 + \frac{(1-p)v_1}{(1-p)v_1 + p v_2} v_2. \tag{16}$$

Both v and $p v_1 + (1 - p)v_2$ are linear combinations of v_1 and v_2. Since $v_2 > v_1$, we must have $\frac{p v_2}{(1-p)v_1 + p v_2} > p$. So $v < p v_1 + (1 - p)v_2$ as v has more weight $\frac{p v_2}{(1-p)v_1 + p v_2} > p$ on the smaller side. By the convexity of the fuel consumption function,

$$m(v) \leqslant m(p v_1 + (1 - p)v_2) \leqslant pm(v_1) + (1 - p)m(v_2). \tag{17}$$

Suppose a vessel is leaving port k and finds that it needs to speed up (shift from v_L to v_H for a fraction $p \in [0, 1]$ of the distance D_k in order to catch up with the schedule. By the above reasoning, even if the speed is continuously adjustable, it is better to choose a constant speed $v = \frac{v_L v_H}{(1-p)v_H + p v_L}$ to achieve the same navigation time. In this case, the fuel consumption will be $D_k m(v)$, which is lower than the fuel consumption $D_k[(1 - p)m(v_L) + pm(v_H)]$ using the two speeds according to (17). However, the exact shape of the function $m(\cdot)$ is required to quantify the amount of reduction.

4.4. Shipping time

We have so far discussed several advantages of the flexible slow steaming strategy. It is obvious that slow steaming will result in longer shipping times than will the fixed high speed strategy. According to formula (1), the scheduled shipping time will be lengthened by $D\frac{v_H - v_L}{v_H v_L}$ if the distance between the origin and the destination is D. The actual shipping time, however, may be shorter because a vessel may speed up for part of the distance. According to formula (14), to maintain a delay probability of less than a certain percentage, at most the same percentage of the distance must be navigated at the high speed. The high speed is adopted only when there is a delay. The flexibility to sail at the high speed mainly helps to reduce the delay probability and the variance of delivery times, but not the shipping time. So our discussion still relies on formula (1) to determine the shipping times.

5. Discussion and numerical experiments

5.1. Impact on liners

Liners normally maintain a fixed frequency of port visits along a service route. The frequency is determined by the projected demand on the service route. Reducing the frequency would mean reducing the liner's own business, since the demand is there. Slowing down the vessels on a service route would mean a longer round-trip time. In order to maintain the same visiting frequency, the liner would have to deploy more vessels on the service route. Take OOCL's Euro-Asia service route via the Suez Canal for example. They used to sail at an average speed of 22.4 knots, resulting in a round-trip time of 56 days. To maintain the service frequency of weekly visits to ports along the route, they needed to deploy eight vessels. After adopting slow steaming with an average speed of 20.2 knots, the round-trip time increased to 63 days. As a consequence, one more vessel was needed to maintain the same service frequency. Although the use of slow steaming required an additional vessel, the total fuel consumption on that service route for the fleet of nine vessels actually decreased by a significant amount. The total round-trip time is

$$T(v) = \frac{D}{v} + \mathbb{E}\left[\sum_{i=1}^{K} W_k + \beta_k \sigma_k\right]. \tag{18}$$

Suppose the liner needs n_L vessels if the flexible slow steaming strategy is adopted and n_H vessels, if fast steaming is adopted to maintain a target service frequency. This means the following relationship should hold for the number of ships n_L and n_H,

$$n_L = \frac{T(v_L)}{T(v_H)} n_H. \tag{19}$$

Note that n_L and n_H can only be integers, so v_L and v_H are in practice not arbitrarily chosen. Liners will normally first determine the service frequency and the number of vessels, and then apply formula (18) to calculate the speed. When slow steaming is considered for a certain service route, the first question is whether to add one or two vessels. Once the number of extra vessels has been determined, the low speed v_L is calculated. The total fuel consumption for n_H vessels on the route traveling at the high speed over a time horizon $[0, T]$ (in practice, T is usually set to a quarter) is $\frac{T}{T(v_H)} n_H Dm(v_H)$. The total fuel consumption for the n_L vessels on the route adopting the flexible slow steaming strategy (i.e., low speed by default and speeding up in case of delay) over the same time horizon is not that easily quantified, but Proposition 3 in Section 4.3 gives an estimate of the upper bound. By Proposition 3, the fuel consumption for n_L vessels using the flexible slow steaming strategy is bounded from the above by $\frac{T}{T(v_L)} n_L [qDm(v_L) + (1-q)Dm(v_H)]$. Relationship (19) and the upper bound in Proposition 3 yield a conservative comparison of the total amount of fuel consumption under two strategies over the same time interval $[0, T]$. The total fuel needed for the fast steaming strategy is $m(v_H)T$ and that for the flexible slow steaming is $[qm(v_L) + (1-q)m(v_H)]T$. Since the fuel consumption function $m(v)$ increases dramatically with speed, flexible slow steaming leads to a big saving of at least

$$\frac{m(v_H) - m(v_L)}{m(v_H)} q100\% \tag{20}$$

on fuel consumption for the whole fleet on a give route.

Of course, bunker cost is only one part of the operating cost. Adding vessels also requires additional crew, so there is an increase in labor cost (also stores, insurances, etc.). But labor cost makes up only a small portion of the operating cost, since only about 20 crew members are needed to operate an 8000-TEU vessel. According to OOCL's data, the saving in bunker cost per vessel per trip as a result of slowing down from 22.4 knots to 20.2 knots is about half a million US dollars, assuming a fuel price of $620/ton, which is significantly larger than the labor cost. Assuming weekly frequency, the total fuel consumption of a ship in a round trip equals the total fuel consumption of all ships in a week. Rough speaking, the annual saving is about 26 million US dollars. This is huge compared with the operating cost (labor, vessel depreciation, etc.) of an additional vessel. For a more detailed discussion on the trade-off between total costs and benefits, see Psaraftis and Kontovas (2010) and Cariou (2011).

Slow steaming requires additional vessels, which are huge investments. But the surplus capacity from heavy investments of past decades and the economic downturn of 2009 have enabled most liners to simply deploy idling vessels to support slow steaming. Once the surplus capacity is absorbed however, liners will need to decide whether to charter or purchase additional vessels. Chartering gives flexibility, but costs more than operating one's own vessel on a per trip basis. Note that the terminal fees determined by the number of visits remains the same since the frequency of visits does not change. All of these costs are important, but they are beyond the scope of this analysis, which focuses on revealing a quantitative relationship between the steaming speed, service quality and fuel consumption. It is then up to the management of the liners to explore the trade-off between better service quality and lower fuel consumption versus longer average shipping times and larger fleet. Building decision models to compare the benefits of chartering and purchasing vessels would be an interesting future research topic.

5.2. Impact on shippers

The major disadvantage of slow steaming for shippers is the longer shipping times. In general, no shipper appreciates longer shipping times. But depending on the businesses they are involved in, shippers may evaluate the trade-off between extended shipping time and improved reliability differently. The complicated Markov chain associated with the slow steaming strategy makes theoretical calculations impossible, but the tradeoffs can be assessed by simulation.

We performed a numerical experiment to see the impact of flexible slow steaming on the variance of the delivery time as well as the delay probability. We simulated one of the service routes of OOCL connecting Asia and the west coast of the U.S. Vessels deployed on this route called at the ports of Shekou (SKZ), Yantian (YAT), Hong Kong (HKG), Los Angeles (LGB), Kao-Hsiung (KHH) and Xiamen (ZIA) before heading back to HKG and finally stopping in SKZ, the home port. To provide a more realistic environment for the simulation, the actual port times for the past three years (provided by OOCL) were used instead of generating random variables with a specified probability distribution. There were about hundreds sample points for each port in the data set. For this particular service route, the round-trip time was about a month and a half. The real distances between ports and the actual steaming speeds were also used. In the simulation, $v_H = 18$ knots and $v_L = 16$ knots. Based on (12), the β's were set to maintain an on-time probability of 60% for each port under fast steaming. Though formula (12) was obtained under the assumption that port times are normally distributed, it worked well for the random data samples too. This actually demonstrated the robustness of the formula in the sense that it gives insights into the system using parameters such as the first two moments which can be estimated from data. The constant fast steaming and the flexible slow steaming strategies were both simulated. Fig. 3 summarizes the simulated delay probability and delivery time variance. The flexible slow steaming strategy improved both metrics, and the improvement became increasingly significant in the course of the voyage. In flexible slow steaming, on average 12.5% of the distance was navigated at the high speed. So the actual fuel saving
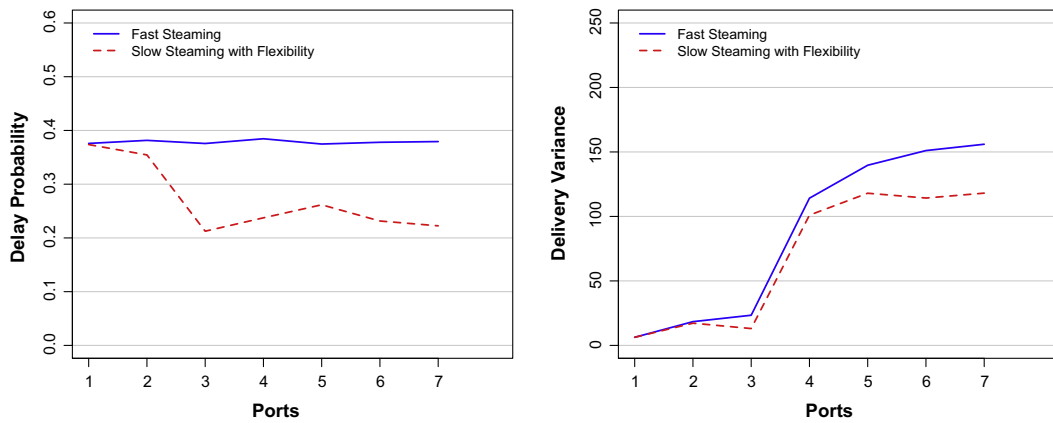
Fig. 3. Delay probability and delivery variance for fast and slow steaming.

compared with fast steaming was $0.875(m(v_H) - m(v_L))D$, which was quite significant. Indeed, this saving was bigger than the theoretic lower bound $0.6(m(v_H) - m(v_L))D$ suggested by Proposition 3.

The effect of the flexible slow steaming strategy is twofold: both the shipping time and the reliability are increased. In practice it is difficult to quantify the positive and negative impact on shippers. After all, different shippers have different preferences for reliability and shipping time. We now briefly illustrate how different strategies affect the safety stock of inventory, since inventory cost could be a key interest of the consignee. Assume that the consignee's demand for the goods being shipped is a random variable with mean λt and variance $\sigma^2 t$ for a period of length t. Let s denote the customer's inventory service level (i.e., the in-stock probability should be larger than s). According to classical inventory theory (c.f. Nahmias, 2009, p. 277), the safety stock is $z_s\sqrt{\sigma^2\mathbb{E}[T] + \lambda^2 Var(T)}$. Here, we focus on a specific customer located at a specific port on a service route, thus the subscript k has been omitted. Fast steaming yields a shorter shipping time than flexible slow steaming, i.e., $\mathbb{E}[T^H] \leqslant \mathbb{E}[T^L]$. But the variance of shipping time is larger for the former strategy than for the latter, i.e., $Var(T^H) \geqslant Var(T^L)$. For example, if we consider a Hong Kong (HKG) based inventory which manages shipment from LBG. Suppose that demand rate λ is scaled to 1, and the demand variance varies from 0.1 to 2. Table 1 presents the safety stock level (based on the above formula) from the simulation in order to maintain the in-stock probability to be at least 95%. We can see that when the demand variance is low, shippers who primarily concern the inventory cost actually benefits from slow steaming. However, shippers facing larger demand variance would prefer fast steaming. We want to point out that there is a "flat" phenomenon: for a wide range of demand variability, the safety stock levels under both steaming strategies do not differ much. This means that slow steaming has relatively little impact for shippers (who are concerned with inventory cost) while it brings huge savings in fuel cost. Customers concerned about inventory cost should estimate the demand rate λ and variance σ^2 to evaluate the trade-off between an increase in shipping time and a reduction in delivery time variability.

Depending on the nature of the products, demand variability can vary. For many consumer goods, the demand variability is quite small, placing them on the low variance side of the above table. For fashion apparel or electronics products, demand variability would be higher. While our table shows the impact of slow and flexible steaming on safety stocks as a function of demand variability, we should note that slow steaming incurs more transit inventory. Since the holding costs of in-transit inventory versus finished goods safety stock can be different, the safety stock effects and transit inventory effects could also be different. Our analysis shows that, in today's ocean transportation industry, flexible slow steaming is a viable alternative strategy, and is suitable for customers who value reliability over short lead times. At the same time this strategy burns much less fuel than a fixed, fast speed strategy, and is therefore good for the environment. The savings resulted from flexible steaming can eventually be divided between shippers and liners. Exactly how it is divided between the two parties can be determined using a pricing mechanism or through contract negotiation. For shippers who are more concerned with inventory costs (both safety stock and transit) and who face high demand variances, they would prefer to have fast delivery, which has been made available by liners charging higher freighter rates as reported by Van de Weijer (2013). The pricing and mechanism design represent another interesting future research topic.

Table 1
Comparison of safety stock level under fast steaming and flexible slow steaming.

Demand variance	0.1	0.2	0.5	0.8	1	1.2	1.5	1.8	2
Fast steaming	20.7	21.7	27.4	35.8	42.1	48.6	58.8	69.3	76.4
Flexible slow steaming	18.7	19.9	26.6	35.8	42.6	49.6	60.5	71.6	79.0

6. Conclusion

We have built a simple yet insightful model to reveal the relationship between delivery time, service quality and bunker cost. Using a comprehensive set of industry data, we obtained the statistical properties of port times and fuel consumption of sea vessels. Service quality can be managed by combining these statistical properties and the model. Our study has quantified the effect of steaming strategies on service quality, bunker cost and shipping time. This can help both shippers and liners to accurately estimate the pros and cons of implementing the slow steaming policy. The insights from this research are summarized as follows:

- Slow steaming with the flexibility to speed up helps to combat the negative effects of the randomness of port times, leading to an improvement in service quality. See Propositions 1 and 2.
- The randomness of port times is usually out of liners' control, though some liners invested in dedicated terminals to gain a degree of control over it. By taking advantage of historical data to identify a pattern of the randomness, and calculating an appropriate buffer time using a simple model, the delay probability can be controlled. The quantitative relationship is given by Eq. (12).
- Calculating an appropriate buffer time for a port also requires incorporating the uncertainties of the preceding ports on a route, as Eq. (12) demonstrates, although the buffer time does not necessarily go up for ports on the latter part of the route.
- An upper bound of the fuel consumption under the flexible slow steaming strategy is presented in Proposition 3, which can help liners control their bunker costs.
- If the speed can be continuously adjusted, liners can further reduce the fuel consumption by choosing an *appropriate* speed according to (16).
- To maintain the same service frequency, the slow steaming strategy requires extra vessels. However, the saving in total fuel consumption is usually higher than the cost of operating the extra vessels. The saving is shown in (20).

This study has also highlighted a few issues that are worth investigating in the future. The first is route design. Liners do adjust their service routes by adding or removing a port or shuffling the sequence of port visits every now and then. But they cannot do it too frequently since a fixed schedule needs to be published to customers months in advance. By exploiting data to characterize the randomness of demand and port times, we can design a better service route. The second issue is dynamic decision making. In this work, we applied a simple policy of speeding up if there is any delay. In reality, we may apply more sophisticated policies. For example, whether to speed up could depend on how many ports remain to be visited on a voyage. It would also be interesting to examine the impact of slow steaming on the design of global supply chains.

Acknowledgements

Appendix A. Technical proofs

Proof of Proposition 1. This result follows quite easily by induction. Since $X_0 = Y_0 = 0, X_1 = Y_1$. So the result is true for $k = 1$. Now, suppose

$$\mathbb{P}(X_l > s) \leqslant \mathbb{P}(Y_l > s), \tag{21}$$

for $l = 1, \ldots, k-1$, we will show that the inequality holds for $l = k$. Note that

$$\mathbb{P}(Y_k > s) = \mathbb{P}(Y_{k-1} + W_k - \mu_k - \beta_k \sigma_k > s) = \int_0^\infty \mathbb{P}(Y_{k-1} > s + \mu_k + \beta_k \sigma_k - \tau) f_k(\tau) d\tau,$$

where f_k is the density function of W_k. The probability for X_k can be computed in a similar way.

$$\mathbb{P}(X_k > s) = \mathbb{P}((X_{k-1} - cD_k)^+ - X_{k-1}^- + W_k - \mu_k - \beta_k \sigma_k > s) = \int_0^\infty \mathbb{P}((X_{k-1} - cD_k)^+ - X_{k-1}^- > s + \mu_k + \beta_k \sigma_k - \tau) f_k(\tau) d\tau.$$

Since $(X_{k-1} - cD_k)^+ - X_{k-1}^- \leqslant X_{k-1}$, we must have for all τ

$$\mathbb{P}((X_{k-1} - cD_k)^+ - X_{k-1}^- > s + \mu_k + \beta_k \sigma_k - \tau) \leqslant \mathbb{P}(X_{k-1} > s + \mu_k + \beta_k \sigma_k - \tau) \leqslant \mathbb{P}(Y_{k-1} > s + \mu_k + \beta_k \sigma_k - \tau).$$

This implies that $\mathbb{P}(X_k > s) \leqslant \mathbb{P}(Y_k > s)$. $\square$

Proof of Proposition 2. Since $S_k(v)$ is deterministic, it is essential only to show that $Var(X_k) \leqslant Var(Y_k)$, for $k = 1, \ldots, K$. We prove this by induction. It is clear that $Var(X_1) = Var(Y_1)$ since $X_1 = Y_1$. Assume that $Var(X_l) \leqslant Var(Y_l)$ for $l = 1, \ldots, k < K$, we need to show that $Var(X_l) \leqslant Var(Y_l)$ for $l = k + 1$. By (20)

$$Var(Y_{k+1}) = Var(Y_k) + Var(W_{k+1}),$$
$$Var(X_{k+1}) = Var((X_k - cD_k)^+ - X_k^-) + Var(W_{k+1}).$$

Thus, it now remains to show that

$$Var((X_k - cD_k)^+ - X_k^-) \leqslant Var(Y_k). \tag{22}$$

Note that

$$(X_k - cD_k)^+ + \min(X_k^+, cD_k) - X_k^- = X_k^+ - X_k^- = X_k.$$

We then have

$$Var(X_k) = Var((X_k - cD_k)^+ - X_k^-) + Var(\min(X_k^+, cD_k)) + 2Cov((X_k - cD_k)^+ - X_k^-, \min(X_k^+, cD_k)).$$

It turns out that the covariance

$$\begin{aligned} Cov((X_k - cD_k)^+, \min(X_k^+, cD_k)) &= \mathbb{E}[(X_k - cD_k)^+ \min(X_k^+, cD_k)] - \mathbb{E}[(X_k - cD_k)^+]\mathbb{E}[\min(X_k^+, cD_k)] \\ &= \mathbb{E}[(X_k - cD_k)cD_k|X_k \geqslant cD_k]\mathbb{P}(X_k \geqslant cD_k) + 0\mathbb{P}(X_k < cD_k) \\ &\quad - \{\mathbb{E}[X_k - cD_k|X_k \geqslant cD_k]\mathbb{P}(X_k \geqslant cD_k) + 0\mathbb{P}(X_k < cD_k)\}\mathbb{E}[\min(X_k^+, cD_k)] \\ &= \mathbb{E}[(X_k - cD_k)|X_k \geqslant cD_k]\mathbb{P}(X_k \geqslant cD_k)(cD_k - \mathbb{E}[\min(X_k^+, cD_k)]) \geq 0 \end{aligned}$$

and the covariance

$$Cov(-X_k^-, \min(X_k^+, cD_k)) = \mathbb{E}(-X_k^- \min(X_k^+, cD_k)) + \mathbb{E}(X_k^-)\mathbb{E}(\min(X_k^+, cD_k)) = \mathbb{E}(X_k^-)\mathbb{E}(\min(X_k^+, cD_k)) \geqslant 0.$$

Thus, we have

$$Var(X_k) \geqslant Var((X_k - cD_k)^+ - X_k^-). \tag{23}$$

So we conclude that the desired result indeed holds. □

References

Agarwal, R., Ergun, O., 2008. Ship scheduling and network design for cargo routing in liner shipping. Transport. Sci. 42 (2), 175–196.
Agarwal, R., Ergun, O., 2010. Network design and allocation mechanisms for carrier alliances in liner shipping. Operat. Res. 58 (6), 1726–1742.
Álvarez, J.F., 2009. Joint routing and deployment of a fleet of container vessels. Marit. Econ. Logist. 11, 186–208.
Bonney, J., 2010. Carriers move full speed into slow steaming. J. Comm.
Brouer, B.D., Alvarez, J.F., Plum, C.E.M., Pisinger, D., Sigurd, M.M., 2014. A base integer programming model and benchmark suite for liner-shipping network design. Transport. Sci. 48 (2), 281–312.
Cariou, P., 2010. Is slow steaming a sustainable mean for reducing liner shipping CO_2 emissions? Euromed Manage. Mare Forum.
Cariou, P., 2011. Is slow steaming a sustainable means of reducing CO_2 emissions from container shipping? Transp. Res. Part D 16, 260–264.
Cariou, P., Cheaitou, A., 2012. The effectiveness of a european speed limit versus an international bunker-levy to reduce CO_2 emissions from container shipping. Transp. Res. Part D 17, 116–123.
Christiansen, M., Fagerholt, K., Ronen, D., 2004. Ship routing and scheduling: status and perspectives. Transport. Sci. 38 (1), 1–18, 2.
Christiansen, M., Fagerholt, K., Nygreen, B., Ronen, D., 2013. Ship routing and scheduling in the new millennium. Eur. J. Oper. Res. 228 (3), 467–483.
Corbett, J., Wang, H., Winebrake, J., 2009. The effectiveness and costs of speed reductions on emissions from international shipping. Transp. Res. Part D 14, 593–598.
Drewry, 2010, 3rd Quarter. Schedule Reliability Insight. Technical report, Drewry Shipping Consultants Ltd.
Fagerholt, K., 1999. Optimal fleet design in a ship routing problem. Int. Trans. Operat. Res. 6 (5), 453–464.
Fagerholt, K., Laporte, G., Norstad, I., 2010. Reducing fuel emissions by optimizing speed on shipping routes. J. Operat. Res. Soc. 61 (3), 523–519.
Feng, X., Mao, B., Feng, X., Feng, J., 2011. Study on the maximum operation speeds of metro trains for energy saving as well as transport efficiency improvement. Energy 36 (11), 6577–6582, 11.
Fransoo, J.C., Lee, C.-Y., 2013. Ocean container transport: an underestimated and critical link in global supply chain performance. Product. Operat. Manage. 22, 253–268.
Kolding, E., 2011. Keynote Speech at TOC Europe. Antwerp, Belgium.
Li, X., Lo, H.K., 2014a. An energy-efficient scheduling and speed control approach for metro rail operations. Transport. Res. Part B: Methodol. 64 (0), 73–89, 6.
Li, X., Lo, H.K., 2014b. Energy minimization in dynamic train scheduling and control for metro rail operations. Transport. Res. Part B: Methodol. 70 (0), 269–284, 12.
Li, C., Qi, X., Lee, C.-Y., in press. How to catch up a delayed journey: Disruption recovery for a vessel in liner shipping. Transport. Sci.
Meng, Q., Wang, S., Andersson, H., Thun, K., 2014. Containership routing and scheduling in liner shipping: overview and future research directions. Transport. Sci. 48 (2), 265–280.
Nahmias, S., 2009. Production and Operations Analysis, sixth ed. Richard D. Irwin Inc., New York.
Notteboom, T.E., 2006. The time factor in liner shipping services. Marit. Econ. Logist. 8, 19–39.
Notteboom, T.E., Cariou, P., 2011. Bunker costs in Container Liner Shipping: Are slow Steaming Practices Reflected in Maritime fuel Surcharges? University Press Antwerp, Antwerp (chapter 11).
Notteboom, T.E., Vernimmen, B., 2008. The effect of high fuel costs on liner service configuration in container shipping. J. Transport. Geogr. 17 (5), 325–337.
Page, P., 2010. Reality shows. J. Comm. 11 (49), 4.
Platt, D.E., Robinson, L.W., Freund, R.B., 1997. Tractable (q, r) heuristic models for constrained service levels. Manage. Sci. 43 (7), 951–965.
Psaraftis, H.N., Kontovas, C.A., 2010. Balancing the economic and environmental performance of maritime transportation. Transp. Res. Part D 15, 458–462.

Psaraftis, H.N., Kontovas, C.A., 2013. Speed models for energy-efficient maritime transportation: a taxonomy and survey. Transp. Res. Part C 26, 331–351.
Qi, X., Song, D.-P., 2012. Minimizing fuel emissions by optimizing vessel schedules in liner shipping with uncertain port times. Transport. Res. Part E: Logist. Transport. Rev. 48 (4), 863–880.
Rana, K., Vickson, R.G., 1991. Routing container ships using Lagrangean relaxation and decomposition. Transport. Sci. 25 (3), 201–214.
Rodrigue, J.-P., Comtois, C., Slack, B., 2009. The Geography of Transport Systems. Routledge, New York.
Ronen, D., 1983. Cargo ships routing and scheduling: survey of models and problems. Eur. J. Oper. Res. 12 (2), 119–126, 2.
Ronen, D., 1993. Ship scheduling: the last decade. Eur. J. Oper. Res. 71 (3), 325–333, 12.
Ronen, D., 2011. The effect of oil price on containership speed and fleet size. J. Operat. Res. Soc. 62, 211–216.
Stahlbock, R., Voß, S., 2008. Operations research at container terminals: a literature update. OR Spectrum 30, 1–52.
Steenken, D., Voß, S., Stahlbock, R., 2004. Container terminal operation and operations research – a classification and literature review. OR Spectrum 26, 3–49.
Ting, S.-C., Tzeng, G.-H., 2003. Ship scheduling and cost analysis for route planning in liner shipping. Marit. Econ. Logist. 5 (4), 378–392.
Van de Weijer, F.A., 2013. Improving Supply Chain Cost and Hinterland Modal Split: An Inventory Theoretic Approach. Master's thesis, Eindhoven University of Technology, Eindhoven, NL.
Vernimmen, B., Dullaert, W., Engelen, S., 2007. Schedule unreliability in liner shipping: origins and consequences for the hinterland supply chain. Marit. Econ. Logist. 9, 193–213.
Wang, S., Meng, Q., 2012a. Liner ship route schedule design with sea contingency time and port time uncertainty. Transp. Res. Part B 46, 615–633.
Wang, S., Meng, Q., 2012b. Sailing speed optimization for container ships in a liner shipping network. Transp. Res. Part E 48, 701–714.
WSC, 2008. Record Fuel Prices Place Stress on Ocean Shipping. Technical report, World Shipping Council.