University of Pennsylvania
ScholarlyCommons

Operations, Information and Decisions Papers

Wharton Faculty Research

2009

Going Bunkers: The Joint Route Selection and Refueling Problem

Omar Besbes
University of Pennsylvania

Sergei Savin

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [Marketing Commons](#), and the [Sales and Merchandising Commons](#)

Recommended Citation

Besbes, O., & Savin, S. (2009). Going Bunkers: The Joint Route Selection and Refueling Problem. *Manufacturing & Service Operations Management*, 11 (4), 694-711. <http://dx.doi.org/10.1287/msom.1080.0249>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/146
For more information, please contact repository@pobox.upenn.edu.

Going Bunkers: The Joint Route Selection and Refueling Problem

Abstract

Managing shipping vessel profitability is a central problem in marine transportation. We consider two commonly used types of vessels—"liners" (ships whose routes are fixed in advance) and "trampers" (ships for which future route components are selected based on available shipping jobs)—and formulate a vessel profit maximization problem as a stochastic dynamic program. For liner vessels, the profit maximization reduces to the problem of minimizing refueling costs over a given route subject to random fuel prices and limited vessel fuel capacity. Under mild assumptions about the stochastic dynamics of fuel prices at different ports, we provide a characterization of the structural properties of the optimal liner refueling policies. For trampers, the vessel profit maximization combines refueling decisions and route selection, which adds a combinatorial aspect to the problem. We characterize the optimal policy in special cases where prices are constant through time and do not differ across ports and prices are constant through time and differ across ports. The structure of the optimal policy in such special cases yields insights on the complexity of the problem and also guides the construction of heuristics for the general problem setting.

Keywords

routing, shipping, refueling, stochastic prices, maritime transportation

Disciplines

Marketing | Sales and Merchandising

Going Bunkers: The Joint Route Selection and Refueling Problem

Omar Besbes*

University of Pennsylvania

Sergei Savin†

Columbia University

November 1, 2008

To appear in *Manufacturing & Service Operations Management*

Abstract

Managing shipping vessel profitability is a central problem in marine transportation. We consider two commonly used types of vessels, “liners” (ships whose routes are fixed in advance) and “trampers” (ships for which future route components are selected based on available shipping jobs) and formulate a vessel profit maximization problem as a stochastic dynamic program. For liner vessels, the profit maximization reduces to the problem of minimizing refueling costs over a given route subject to random fuel prices and limited vessel fuel capacity. Under mild assumptions about the stochastic dynamics of fuel prices at different ports, we provide a characterization of the structural properties of the optimal liner refueling policies. For trampers, the vessel profit maximization combines refueling decisions and route selection which adds a combinatorial aspect to the problem. We characterize the optimal policy in special cases where: *i*) prices are constant through time and do not differ across ports, and *ii*) prices are constant through time and differ across ports. The structure of the optimal policy in such special cases yields insights on the complexity of the problem and also guides the construction of heuristics for the general problem setting.

1 Introduction

Water transportation is an important component of the U.S. transportation industry, accounting for more than \$23 billion dollars in revenues according to the U.S. Economic Census (2002). The Marine Transportation System National Advisory Council (2000) estimates that the marine import-export trade alone accounts for nearly 7% of the U.S. gross domestic product. In recent years, increased competition and global downturn in the shipping industry have been putting downward pressure on the revenues of shipping companies, while increased safety regulations and fuel prices continued

*The Wharton School, e-mail: obesbes@wharton.upenn.edu

†Graduate School of Business, e-mail: svs30@columbia.edu

to increase companies' operating costs. The changing economic environment has prompted many companies to abandon a "status quo" complacency and search for new ways to maintain and optimize their profitability.

A typical marine transportation company operates a fleet of ships which belong to one of two broad groups. *Liners* are vessels that follow the same cyclical route comprised of a string of ports, while *trampers* are vessels-for-hire for which the next destination is selected according to the set of available transportation jobs. The fundamental "routing" distinction between these two groups of vessels translates into important differences in how liners and trampers are managed. A liner brings in a steady stream of revenues, and a major managerial challenge is to minimize its refueling costs given a limited vessel fuel capacity and unknown future fuel prices. A tramper, on the other hand, has a choice of ports it visits in the future, and, thus, can control both its revenues and refueling costs. Both types of vessels use "bunkers," a by-product of the oil refining process, as a fuel - and the bunkers' prices are highly unpredictable and can exhibit significant variation across the ports on the same date. Note that vessels typically also use diesel oil for auxiliary power and while such a component could be included in the model, we do not consider it to keep the exposition clear.

In the present paper, we propose a model for bunkers price dynamics in conjunction with a novel model which describes a single-vessel profit optimization problem. In the case of a liner, the profit optimization reduces to minimizing the refueling costs. We formulate this problem as a long term average stochastic dynamic program where the decisions are the refueling amounts at each port the vessel visits on its fixed route. In the presence of random fuel prices, this problem becomes a variant of the stochastic capacitated inventory management problem, whose solution is shown to be a capacity-adjusted state-dependent buy-up-to policy. For a tramper, the refueling dynamic programming problem is blended with combinatorial optimization of the vessel's route. For a static (deterministic or stochastic) routing policy, the tramper problem reduces to a liner instance, but the addition of dynamic routing significantly complicates the analysis of optimal profit management policies. The fact that refueling and routing decisions are interconnected makes the tramper problem unique since it cannot be reduced to classes of problems analyzed in the stochastic routing literature (we return to this point in the next section). In particular, in general a tramper cannot ignore bunkers prices when selecting a route to maximize profitability significant as price differences among ports located in the same region might exist. For example, on November 30, 2007, the following bunkers prices (per metric ton) were recorded for a few ports in North and Central America: \$473 (Philadelphia), \$445 (Houston), \$521 (Los Angeles), \$460 (Panama).

The main contributions of the present work can be summarized as follows:

1. We develop a new modeling framework to analyze a single-vessel profit optimization problem as a blend of combinatorial route selection and dynamic programming refueling problems in

the presence of stochastic revenues and fuel prices.

2. For vessels which follow a static routing policy (liners), we show that the optimal refueling policy is of the buy-up-to form (Proposition 1) and that the value of the buy-up-to level necessarily belongs to a finite, potentially small set. Under additional assumptions on the stochastic monotonicity of the underlying fuel price processes, we also establish monotonicity properties of the optimal buy-up-to levels (Proposition 2).
3. For vessels which combine refueling with route selection (trampers), we demonstrate that, in the general case, while the optimal refueling quantities need not be monotone in on-board inventory, the optimal post-refueling inventory levels are (Proposition 3). In addition, we investigate several special cases of the tramper problem. In particular, in the case when the bunkers prices are constant across time and equal across ports, we show that the problem can be reduced to one of finding the best profit-to-time ratio cycle on a network of ports. Also, in the case when the bunkers prices are constant across time but *differ* across ports, we show that the solution forms a cycle in a generalized location-inventory space (Proposition 4). In addition, we investigate the impact of vessel capacity (Proposition 5) and price stochasticity on optimal routing policies.
4. Based on the above results, we develop heuristics for the general case of the tramper problem and derive performance bounds for those (Propositions 6 and 7).

Our analysis of the vessel profit optimization problem includes further modeling of the stochastic fuel prices in which, at each port, the influence of global oil price dynamics is augmented by the contribution of the local Markovian demand-supply dynamics.

Our paper is organized as follows. The next section provides a review of related work. Section 3 introduces our vessel profit maximization model. Section 4 is devoted to modeling of the stochastic dynamics of bunkers prices. The special case of liners is analyzed in Section 5, while Section 6 focuses on trampers. We conclude by discussing possible future research directions for our work.

2 Review of Related Work

Traditional marine fleet management research literature outlines a hierarchy of vessel management problems (Christiansen et al. 2004, Christiansen et al. 2005). At the top of the hierarchy are the models which study strategic, long-term decisions, such as fleet sizing and route design (Dantzig and Fulkerson 1954, Richetta et al. 1997, Imai and Rivera 2001, Cho and Perakis 1996). On a more tactical level, fleet management models focus on routing and scheduling of a fixed number of vessels. In particular, existing papers on liner operations exclusively focus on deterministic models

of fleet deployment and cargo booking (Perakis and Jaramillo 1991, Powell and Perakis 1997). In the context of airlines, linear programming formulations for fuel management have been proposed in the literature (where prices are taken to be constant over time). See, e.g., Darnell and Loffin (1977) and Stroup and Wollmer (1992) and references therein. These studies relate to the liner study presented here, as the route of the flights is fixed. While it is natural to take prices to be constant over the horizon of a route given the short time elapsed between each stop in the context of airlines, stochasticity of prices plays a more important role in the context of marine shipping and is explicitly taken into account in our liner formulation.

In its most general form, profit maximization at the vessel level represents a blend of combinatorial optimization (route selection) and stochastic dynamic programming (minimization of refueling costs).

The problem of controlling refueling costs is related to a number of papers on inventory management in the presence of random prices. In the first paper in this area, Kalymon (1971) generalizes the classical inventory model by Scarf (1960) (formulated for the setting with convex holding and shortage costs and a setup cost) by allowing product prices to follow a non-stationary Markov process. In the finite-horizon setting, the optimal replenishment policy in period n is shown to be of the $(s_n(p), S_n(p))$ form, where p is the realized product price in that period. Song and Zipkin (1993) consider a continuous time infinite-horizon inventory model where the product demand rate varies in a Markov-modulated fashion with an underlying “state-of-the-world” variable: if the “world” is in state i , the demand is assumed to follow a Poisson process with rate λ_i . Under the assumptions of full demand backlogging, stochastic leadtimes, and convex ordering and holding costs the optimality of a state-dependent base-stock policy (without the order setup cost) and of $(s_n(i), S_n(i))$ policy (with the setup cost) is established. This framework was extended by Ozekici and Parlar (1999) to include the influence of the Markovian random environmental process on the demand, supply and cost parameters. An interesting variant of the random-price setting is analyzed in Moynadeh (1997) where a continuous-time infinite-horizon setting is used to study an inventory problem with fixed price discounts offered at random times. The price-dependent two-bin (R, s, Q) replenishment and stocking policy is employed: when the deal is offered, an $(s, s + Q)$ replenishment rule is used, while when the inventory is completely depleted, the order of $R \leq s + Q$ units is placed. While for a fixed vessel route, the refueling cost minimization problem resembles classical inventory replenishment problems with random prices and deterministic demand values, two features distinguish it from the variants studied in the literature: the finite fuel capacity of a vessel and the availability of the entire set of prices at all ports at any point in time. In other words, we deal with a capacitated inventory problem for which the state information is a vector of fuel prices which includes the price “now” as well as indicators for future prices (at all ports in the network). The presence of geographical dependence of prices is an aspect that is specific to

the problem we analyze and has not, to the best of our knowledge, been studied in the inventory management literature.

The stochastic nature of future potential rewards places route selection into the same class of problems as, for example, the stochastic vehicle routing problem. However, an important distinction of the routing aspect of the problem under consideration from the classical stochastic vehicle routing literature resides in the the formulation itself. In our setting, the objective of the decision-maker is to maximize the long-term average profits with constraints coming into play only through the refueling decisions. In contrast, the typical stochastic vehicle routing problem consists of minimizing the total cost of a travel plan such that *each* node in the network is visited at least once, with potential side constraints. We refer the reader to Gendreau et al. (1995) for a review of stochastic vehicle routing.

An important observation is that the tramper problem, in the absence of refueling and in the special case of deterministic rewards, actually reduces to the problem of finding the cycle with highest revenue per unit of time over a network, which was studied in Dantzig et al. (1956). We build on this connection to understand the structure of optimal policies in the presence of refueling decisions.

Finally, we mention that our price model describes the evolution of bunkers prices in terms of the dynamics of two independent Markov chains. Our approach to modeling bunkers prices is related in spirit to the work of Hamilton (1989) which incorporates Markovian “regime changes” into the dynamics of energy prices. We also refer the reader to Hamilton and Susmel (1994), Either and Mount (1998) and Noel (2007a,b) for more recent applications of this approach.

3 Optimizing Vessel Profits: The Model

We consider a vessel with fuel capacity C (measured, e.g., in metric tons (mts)) traveling through a network of N nodes, each representing a port. Since a day is a natural time unit in our setting, we use a discrete-time formulation for the vessel profit optimization problem. Each node in the network of ports is assumed to be connected to any other node by an arc. We use $\mathcal{N} = \{1, 2, \dots, N\}$ to denote the set of all nodes, and $\mathcal{A} = \{(i, j) : i \in \mathcal{N}, j \in \mathcal{N}\}$ to denote the set of arcs connecting the nodes in $\mathcal{N}$. On the complete graph $(\mathcal{N}, \mathcal{A})$ we define the following set of measures:

1. $\mathbf{d} = \{d_{ij} > 0 | (i, j) \in \mathcal{A}\}$ and $\boldsymbol{\tau} = \{\tau_{ij} \in \mathbb{Z}_+ | (i, j) \in \mathcal{A}\}$ are the amount of fuel (e.g., in mts) and the time (in days), respectively, required to travel through the arcs of $\mathcal{A}$.
2. For each node $i \in \mathcal{N}$, the set $O(i) = \{j : (i, j) \in \mathcal{A}, d_{ij} \leq C\}$ denotes nodes which are “reachable” from i without stopping at any other port. Without loss of generality, we assume

that the capacity of the vessel is such that any two nodes $i \in \mathcal{N}$ and $j \in \mathcal{N} \setminus O(i)$ are connected by a finite set of arcs, i.e., that there exists a finite set of nodes $i_1, \dots, i_M$ such that $i_1 \in O(i), i_2 \in O(i_1), \dots, j \in O(i_M)$ for some $M \geq 1$.

3. For each time period $t = 1, 2, \dots$, the set of rewards associated with shipping contracts on every arc $(i, j) \in \mathcal{A}$ is denoted as $\mathbf{r}(t) = \{r_{ij}(t) | (i, j) \in \mathcal{A}\}$. The reward $r_{ij}(t)$ associated with traveling on a given arc $(i, j) \in \mathcal{A}$ at time period t is assumed to be a realization of a non-negative bounded discrete i.i.d. random variable with a static probability distribution $p_{ij}(r) = \text{Prob}(r_{ij}(t) = r)$. In our analysis, we assume that each shipping contract is associated with a single arc (i, j) - an assumption valid for a large majority of contracts we observed in practice.
4. For each time period $t = 1, 2, \dots$, the set of fuel prices for all ports is denoted as $\mathbf{P}(t) = \{P_i(t) | i \in \mathcal{N}\}$. We assume that for any time period t and for any port $i \in \mathcal{N}$ the price $P_i(t)$ can only take one of the L values $\{P^{(1)}, \dots, P^{(L)}\}$. In addition, we assume that the vector of prices $\mathbf{P}(t)$ follows a Markov process such that any price state can be reached from any other price state with a positive probability in a finite number of transitions.

Consider a vessel arriving at time t to port i with fuel inventory I and observing the values of fuel prices $\mathbf{P}$ and rewards $\mathbf{r}$. Let a stationary unichain policy $\pi : (i, I, \mathbf{P}, \mathbf{r}) \mapsto (j, q)$ denote the choice of the next port to visit $j \in \mathcal{N} \setminus \{i\}$ ¹ and the refueling amount $q \in Q(i, j, I) := \{q | q \geq 0, d_{ij} \leq I + q \leq C\}$. Let $V_k^\pi(i, I, \mathbf{P}, \mathbf{r})$ be the total expected profit earned in the next k port-to-port transitions under a stationary policy π , and let $t_k^\pi(i, I, \mathbf{P}, \mathbf{r})$ be the expected time associated with these k transitions. In our analysis we use the long-run expected profit per unit of time

$$\lambda = \lim_{k \rightarrow \infty} \left(\frac{V_k^\pi(i, I, \mathbf{P}, \mathbf{r})}{t_k^\pi(i, I, \mathbf{P}, \mathbf{r})} \right) \quad (1)$$

as an optimization criterion for the vessel. Note that since we assume the travel times between ports to be deterministic, the analysis we conduct uses the standard approaches for semi-Markov decision processes as outlined in Bertsekas (2000). Clearly, under the assumptions on the reachability of nodes and on the reward and price dynamics presented above, the limit in (1) is well-defined, and its value is independent of the initial state. Following Bertsekas (2000), we can express the Bellman equation for λ as

$$\begin{aligned} & h(i, I, \mathbf{P}, \mathbf{r}) \\ &= \max_{j \in O(i)} \left(r_{ij} - \lambda \tau_{ij} + \max_{q \in Q(i, j, I)} \left(-P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}, \mathbf{r}'} \left[h(j, I + q - d_{ij}, \mathbf{P}', \mathbf{r}') \right] \right) \right), \end{aligned} \quad (2)$$

¹Note that here, we assume that the ship never idles at a port. While this is typical of most practical settings, allowing for idling is possible at the expense of additional notation.

where $h(i, I, \mathbf{P}, \mathbf{r})$ is a function defined on the state space of the problem with $h(i_0, I_0, \mathbf{P}_0, \mathbf{r}_0) = 0$ for some (arbitrarily chosen) state $(i_0, I_0, \mathbf{P}_0, \mathbf{r}_0)$, and $\mathbf{P}'|\mathbf{P}$ denotes that the expectation with respect to $\mathbf{P}'$ (the vector of prices observed τ_{ij} units of time after $\mathbf{P}$ is observed) taken conditional on the value of $\mathbf{P}$.

The outer maximization operator in (2) represents the routing decision, while the inner maximization operator corresponds to the refueling decision. For simplicity, we assume that the refueling is instantaneous.

The formulation as well as the analytical results to be presented are valid for *any* Markovian model for the fuel price dynamics, except when explicitly stated otherwise. In the next section, to finalize the specification of the model, we develop a particular fuel price model and estimate its parameters using actual pricing data for 18 ports for the period of January-June 2005. This model will then be used in our numerical experiments. We return to the analysis of the liner and tramper's problems in Sections 5 and 6, respectively.

4 Modeling Stochastic Dynamics of Bunkers Fuel Prices

In practice, selecting an appropriate price dynamics model is a challenging task due to strong time and geographical correlations in bunkers prices. In our analysis, we have selected the general form of the Markov process $\{\mathbf{P}(t) : t \geq 0\}$ based on the following two fundamental features of the bunkers price dynamics identified by practitioners. On the one hand, the main, global driver of the bunkers fuel prices on a given day is crude oil price². This dependence stems from the fact that the bunkers fuel is a by-product of the oil refining process. In addition to this global effect, on the level of a particular port, the bunkers price is also influenced by the interactions between supply of and demand for the fuel. This local supply effect is believed to be weaker than the global crude oil price effect.

Below we propose a model of bunkers price dynamics which incorporates the two effects just mentioned. Let $P_0(t)$ denote the spot price of crude oil (in \$ per barrel) in period t . Then, the bunkers fuel price $P_i(t)$ (in \$ per metric ton) at location i in period t is described as follows:

$$P_i(t) = \gamma P_0(t) + \alpha_i(t) + \epsilon_i(t), \quad (3)$$

where γ is a constant, $\alpha_i(t)$ are Markov processes independent of P_0 and independent across ports, each with m possible states, and $\epsilon_i(t)$ is a stationary random variable with $\mathbb{E}[\epsilon_i(t)] = 0$, $\mathbb{E}[\epsilon_i^2(t)] = \sigma_i^2 < \infty$ for all i and t . We also assume that all $\epsilon_i(t)$ are independent across time

²The two main types of crude oil which are typically used by traders as indicators for the bunkers prices are Brent (traded on London's Intercontinental Exchange) and WTI (Western Texas Intermediate, traded on the New York Mercantile Exchange).

and across ports (and of P_0 and α_i 's). Introducing a dependence on oil price in model (3) follows the established pricing models for other products of oil refining process, such as gasoline or home heating oil (Borenstein *et al.* 1997, Asche *et al.* 2003, Kaufmann and Laskowski 2005); the introduction of location-dependent terms α_i and ϵ_i are important additions. The interpretation of the model (3) is as follows: γ represents a global price-adjustment factor, α_i are local (geographical) supply correction factors for the bunkers prices and $\epsilon_i(t)$ are geographical adjustments due to other factors. The role of the oil price, and, thus, the presence of γ in (3) is straightforward in models of any oil-based products. The necessity to use $\epsilon_i(t)$ term stems from the presence of unexpected and potentially *short-lived* local factors such as sudden surge of the demand for bunkers due to unanticipated arrival of several tramper vessels at a particular port. On the other hand, the α_i terms stand to represent the potentially *longer-term* effects on the bunkers price stemming from juxtaposition of global changes in the state of oil-refining industry and local demand and supply conditions (the evolution of the α_i 's is related to the “regime changes” of Hamilton (1989)). At any port, we describe the Markovian dynamics of the local supply correction factors using a simple 2-state form: α_i is assumed to take “high” values when the fuel supply at the port is constrained and “low” otherwise; i.e. α_i , $i = 1, \dots, N$ can have two possible states: $\alpha_i^\pm = A_i \pm \delta_i$, where A_i plays the role of the average price premium/discount over the value of $\gamma P_0(t)$ associated with port i , and δ_i characterizes the amplitude of the price response at port i to the changes in the local demand-supply balance. The transitions between these two supply states are described at port i by a local “inertia” parameter η_i , which stands for the probability for the local supply state to remain unchanged on the next day. Thus, if the port i is in state $A_i + \delta_i$ on some day, it will remain in the state $A_i + \delta_i$ on the next day with probability η_i and transfer to the state $A_i - \delta_i$ with probability $1 - \eta_i$. We note that according to (3) the state of the price system in period t is described by $2N + 1$ values: the oil price $P_0(t)$, the set of local supply corrections α , and the perturbation vector ϵ . The motivation for including the “inertia” terms into the pricing model (3) is illustrated in Figure 1 which plots actual bunkers prices observed in 3 large ports, Rotterdam, Houston and Singapore, during the period January-June 2005 (the total of 98 bunkers trading days).

We observe that, in addition to the general common price trend and to the random daily shocks, the bunkers prices for these 3 ports exhibit an absolute ranking with some longer-term stability. For example, for the first 10 or so trading days, the price ranking of 3 ports, from highest to lowest, is Houston-Singapore-Rotterdam. After that, the ranking becomes, almost uninterruptedly for the next 30 days, Singapore-Houston-Rotterdam. For the next 5 or so days, the ranking is Singapore-Rotterdam-Houston. After that, almost uninterruptedly, the Singapore-Houston-Rotterdam ranking is restored. The “inertial” Markovian structure of the local supply correction terms α_i helps to model such behavior.

Bunkers pricing data is not publicly available and shipping companies resort to specialized firms

Figure 1: Bunker prices in 3 world’s largest ports, January-June 2005 (98 bunkers fuel trading days).

for historical worldwide pricing information. We obtained pricing information from such a firm in the form of actual prices for $N = 18$ ports for the period of January-June 2005 ($k = 98$ trading days). A spreadsheet with price information is enclosed as a part of an online appendix.

The details of the parameter estimation procedure for model (3) are provided in Appendix B. The value of the parameter γ representing a conversion factor between the time average of the oil price expressed in \$/bbl and the time-and-location average of bunkers prices expressed in \$/mts is estimated as 4.43, and the values of other parameter estimates for three of the world largest ports Rotterdam, Houston, and Singapore, as well as some descriptive statistics for all ports are presented in Table 1. Note that, as Table 1 indicates, the expected price premia/discounts A_i vary substantially over the geographical locations: from a deep expected discount of \$31.6 (observed at the port of Antwerp, Belgium), to a heavy premium of \$34.0 (observed at the port of Yokohama, Japan). Yet, these expected values do not tell the entire story: the amplitudes of the Markovian local supply corrections are around \$30, decisively influencing the resulting price ranking of ports. Finally, very high estimates for η_i indicating extreme “stickiness” of the local supply corrections are largely due to a 2-state nature of the Markovian supply correction model and high values of the amplitudes of supply corrections: note that the transition between $A_i + \delta_i$ and $A_i - \delta_i$ prices involves a price change of $2\delta_i$, or, on average, of \$57. Clearly, a reasonable model should predict that such dramatic price changes happen very infrequently in order to produce reasonable values

for the average price change due to “jumps” between local supply states. For example, the average “inertia” coefficient equal to 0.98 and the average δ equal to \$28.6 correspond to equivalent daily price change of $(1 - 0.98) \times 2 \times \$28.6 = \$1.04$. An increase in a number of modeled possible states for the local bunkers supply chain will result in a reduction in the estimates for the “inertia” coefficients associated with each state.

Parameter	Rotterdam	Houston	Singapore	Max (all ports)	Min (all ports)	Mean (all ports)
A , \$/mts	−28.4	−13.4	−0.42	34.0	−31.6	0
δ , \$/mts	29.7	31.7	30.8	31.7	24.6	28.6
η	0.99	0.93	0.99	0.99	0.93	0.98
σ , \$/mts	14.8	13.9	14.2	18.1	10.9	14.2
$\frac{\sigma}{\delta}$	0.50	0.44	0.46	0.61	0.38	0.50

Table 1: Parameters of the bunkers fuel price dynamics.

In order to fully characterize the proposed model (3), we turn attention to the last remaining element of the stochastic dynamics of bunkers prices, namely, the stochastic dynamics of crude oil prices. The earlier literature on the subject (Brennan and Schwartz 1985, Paddock et al. 1988, Smith and McCardle 1998) has established the geometric Brownian motion (GBM), in both continuous and discretized versions, as a standard for the description of the crude-oil dynamics. In the discretized version of the GBM model the relative changes in the oil prices for any two consecutive days are assumed to be i.i.d. normal random variables:

$$\frac{P_0(t+1) - P_0(t)}{P_0(t)} \sim N(\mu_0, \sigma_0). \quad (4)$$

More recently, however, a number of studies (Laughton and Jacoby 1995, Cortazar and Schwartz 1994, Dixit and Pindyck 1994, Smith and McCardle 1999) have argued that the unlimited asymptotic variance associated with the GBM model may not be a good descriptor of actual price dynamics. As an alternative, these studies suggested to use mean-reverting processes, such as an $AR(1)$ process

$$P_0(t+1) = \lambda P_0(t) + (1 - \lambda)\hat{P}_0 + \varepsilon_0(t), \quad (5)$$

where $\hat{P}_0$ is the long-term price level to which oil prices “revert”, and $\varepsilon_0(t)$, $t = 0, 1, \dots$ are i.i.d. normal random variables. By analyzing the actual crude oil data for the period January-June 2005 (see the estimation details in Appendix B), we have found that the mean-reverting model provides a better description of the crude oil price dynamics. It is this model that we use in our numerical studies.

5 The Case of a Liner

The fixed-route feature of a liner trajectory simplifies the analysis of the vessel profit management problem (2) and allows one to characterize optimal refueling policies.

5.1 Structural Properties of the Optimal Refueling Policies

For a liner, at each port i the destination $j \in O(i)$ is known and fixed. We let $j^{(0)}(i) = i$ and $j^{(k)}(i)$ denote the k^{th} port visited when starting at port i ³. With a slight abuse of notation, we let $j(i)$ denote $j^{(1)}(i)$. Similarly, we let $d^{(k)}(i)$ denote the consumption needed for the next k ports, i.e.,

$$d^{(k)}(i) = \sum_{m=0}^{k-1} d_{j^{(m)}(i)j^{(m+1)}(i)}.$$

When the route is fixed, the revenue stream is known and is unaffected by the refueling decisions of the firm. In this context, the objective of the firm is to minimize the long run average fuel costs along the sequence of ports visited. If we let μ denote the optimal cost per unit of time, the Bellman equation can be written as follows

$$\begin{aligned} w(i, I, \mathbf{P}) \\ = -\mu \tau_{ij(i)} + \min_{q \in Q(i, j(i), I)} \left(P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}} \left[w(j(i), I + q - d_{ij(i)}, \mathbf{P}') \right] \right), \end{aligned} \quad (6)$$

where $w(i, I, \mathbf{P})$ is a function defined on the state space of the problem with $w(i_0, I_0, \mathbf{P}_0) = 0$ for some (arbitrarily chosen) state $(i_0, I_0, \mathbf{P}_0)$.

The following result summarizes structural properties of the optimal cost function and the optimal refueling policy:

Proposition 1 *a) the function $w(i, I, \mathbf{P})$ associated with the optimal cost is a convex function of fuel inventory I ,*

b) there exists a price-dependent function $S(i, \mathbf{P})$ taking values in $[d_{i,j(i)}, C]$ such that an optimal refueling policy is given by:

$$Q^*(i, I, \mathbf{P}) = (S(i, \mathbf{P}) - I)^+ \quad (7)$$

c) the function $S(i, \mathbf{P})$ takes values in $B(i)$, where $B(i) = \{d^{(l)}(i) : l = 1, 2, \dots, l^(i)\} \cup \{C\}$, with $l^*(i) = \max\{k : d^{(k)}(i) \leq C\}$.*

Proposition 1b) describes a capacitated version of the base-stock inventory policy and admits a straightforward interpretation. On the one hand, when refueling is necessary at port i (corresponding to $I \leq d_{ij^*(i)}$), it should be performed up to the level equal to $S(i, \mathbf{P}(t))$. On the other hand,

³Note that as stated, any liner has to travel through a cycle in the network of ports. The analysis easily extends to a case where the liner follows any repeated route (not necessarily a cycle in the network of ports) at the expense of additional notation.

if there is enough fuel to get to the next port $j^*(i)$, i.e. $I > d_{ij^*(i)}$, the refueling at i should be done if and only if $S(i, \mathbf{P}(t)) > I$. By Proposition 1c), the optimal buy-up-to level belongs to the set $B(i)$; in other words, whenever bunkers are purchased, they are purchased to either arrive at one of the next ports on the route with exactly zero inventory, or they are purchased to fill up the tank up to capacity. This property allows one to restrict the set of possible values for $S(i, \mathbf{P})$ when searching for an optimal policy.

Proposition 1 describes the structure of the optimal refueling policy for any general Markovian price process $\{\mathbf{P}(t) : t \geq 0\}$. Sharper characterization of the properties of the capacity-adjusted fuel-up-to levels $S(i, \mathbf{P}(t))$ is possible for specific models of the Markovian price dynamics. For example, based on the structure of the fuel pricing dynamics described in (3) in Section 4, one can analyze the properties of the fuel-up-to-levels $S(i, P_0, \epsilon, \alpha)$ (note that since the values of the local supply correction terms α_i are assumed to be observable, expressing the state of the system in terms of the prices $\mathbf{P}$ is equivalent to specifying the values of corrections ϵ). In particular, let $F(P_0(t+1)|P_0(t))$ denote the CDF of the oil price in period $t+1$ given that the oil price in period t is $P_0(t)$. Consider the following assumptions on the shape of F :

Assumption 1 (Crude oil price dynamics)

- i) $F(P_0(t+1)|P_0(t))$ is a non-increasing function of $P_0(t)$ for any $P_0(t+1)$.
- ii) $P_0(t) - \mathbb{E}(P_0(t+t')|P_0(t))$ is a non-decreasing function of $P_0(t)$ for all $t' \geq 0$.

Assumption 1 can be rationalized as follows. Let $P_0(t+1)|P_0(t)$ be the (random) oil price in period $t+1$ given that in period t the oil price is $P_0(t)$. Then, Assumption 1 (i) states that $P_0(t+1)|P_0(t) = P_H$ stochastically dominates $P_0(t+1)|P_0(t) = P_L$ for all $P_L \leq P_H$. In particular, it implies that the conditional expectation $\mathbb{E}(P_0(t+1)|P_0(t))$ is a non-decreasing function of $P_0(t)$. At the same time, Assumption 1 (ii) limits the rate at which this conditional expectation grows as a function of $P_0(t)$. Note that both the mean-reverting model as well as GBM model with non-positive drift satisfy this Assumption. Assumption 1 serves as a sufficient condition for the monotonicity of the fuel-up-to levels $S(i, P_0, \epsilon, \alpha)$:

Proposition 2 *Suppose that bunkers fuel prices evolve according to (3). Then the fuel-up-to level $S(i, P_0, \epsilon, \alpha)$ is a non-increasing function of ϵ_i (or equivalently P_i for fixed P_0 and α). If, in addition, Assumption 1 holds, then the fuel-up-to level $S(i, P_0, \epsilon, \alpha)$ is a non-increasing function of P_0 for fixed ϵ and α .*

In particular, we note that under Assumption 1, when the local supply corrections are constant across time, the values of the α_i 's do not influence the monotonicity of the optimal fuel-up-to levels. However, when Assumption 1 is not satisfied, it is easy to find examples where the monotonicity of the fuel-up-to levels as functions of the oil price P_0 breaks down.

In Appendix C, we illustrate numerically the properties of the optimal fuel-up-to levels described in

Propositions 1 and 2 as well as the performance of simple heuristics, which require more moderate computational effort.

6 The Case of a Trampler

In this section we consider a profit maximization problem for a vessel whose route can be dynamically adjusted based on the set of available shipping jobs and the set of observed fuel prices. In particular, we focus on the case of a so-called “spot trampler”, for which the set of available jobs becomes known only after the completion of a previous shipping job, upon arrival at a discharge port, and which reflects the practice of several trampler companies we collaborated with.

The case of a trampler vessel is far more complex and less amenable to analysis due to the additional routing decision that needs to be made at every port. Section 6.1 illustrates some of the complications that arise in the context of the trampler problem when contrasting it to the liner. Section 6.2 analyzes the trampler problem in various special cases where either the optimal solution can be characterized or simple policies can be shown to be near-optimal. Based on the latter analysis, Section 6.3 develops potential heuristics for the general trampler problem and illustrates their performance.

6.1 Optimal Refueling Decisions: Monotonicity Properties

Below, we re-express the Bellman equation given in (2) as

$$h(i, I, \mathbf{P}, \mathbf{r}) = \max_{j \in O(i)} (r_{ij} - \lambda \tau_{ij} + D_{ij}(I, \mathbf{P})), \quad (8)$$

where

$$D_{ij}(I, \mathbf{P}) = \max_{q \in Q(i, j, I)} \left(-P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}, \mathbf{r}'} \left[h(j, I + q - d_{ij}, \mathbf{P}', \mathbf{r}') \right] \right) \quad (9)$$

with $Q(i, j, I) = \{q | q \geq 0, d_{ij} \leq I + q \leq C\}$. Note that while concavity of h with respect to inventory is preserved under the transformation on the right-hand side of (9), the combinatorial nature of the route selection problem in (8) breaks this concavity down. This, in turn, affects the monotonicity properties of the optimal refueling policies.

One can derive sensitivity properties for the optimal refueling values. In particular, assume that while in port i , node j is selected as the next destination, and let

$$q^*(i, j, I, \mathbf{P}) = \min \left(\arg \max_{q \in Q(i, j, I)} \left(-P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}, \mathbf{r}'} \left[h(j, I + q - d_{ij}, \mathbf{P}', \mathbf{r}') \right] \right) \right)$$

be the smallest optimal refueling quantity in this case. Then,

Proposition 3 *$I + q^*(i, j, I, \mathbf{P})$ is a non-decreasing function of I for any $i = 1, \dots, N$, $j \in O(i)$, and for any current price vector $\mathbf{P}$.*

While the monotonicity of the fuel-up-to level with respect to the on-board inventory is to be expected, it may be possible for the refueling quantity itself to be non-monotone due to the influence that the routing decisions exert on the refueling process.

6.2 Optimal and Near-Optimal Solution Structure: Special Cases

An important feature of the tramper problem is, in general, a strong interaction between the routing and the refueling decisions. Indeed, the latter decisions need not decouple in the general case and the fact that they are selected jointly presents unique challenges. In this section, we develop an understanding of the solution structure in various special cases. This will highlight how various features of the tramper problem, independently of each other, create complex interactions between refueling and routing decisions. We focus on three key features associated with the tramper problem: *i*) price heterogeneity (across ports); *ii*) price stochasticity; and *iii*) vessel capacity constraint. We will illustrate how each of these features introduces strong interactions between the refueling and routing decisions. In addition to these insights, the analysis of the special cases will also be the basis for the development of heuristics with reduced computational requirements in Section 6.3.

In order to analyze the various impacts of the model features above, we will start from a base case where all the features mentioned above are absent and then study the impact of adding one of those features. In particular, our base case will be a variant of the vessel profit management problem (2) with deterministic reward values $\mathbf{r}$, deterministic, constant (across time) and uniform (across ports) fuel prices, i.e., $P_i = P_c, i \in \mathcal{N}$ for some common value P_c . In addition, we assume that the consumption between ports and corresponding travel times are proportional, i.e., that for all $(i, j) \in \mathcal{A}$, $d_{ij} = \nu \tau_{ij}$ for some $\nu > 0$.

6.2.1 The Base Case: Deterministic Uniform Fuel Prices

In the base case, prices are equal across all ports and constant, and hence, without loss of generality, one can restrict attention to policies where the firm replenishes the exact amount required to reach the port it selects to travel to next. Under any such policy, the fuel on board upon arrival at a port will always be exactly zero and it is possible to associate a profit with each arc $(i, j) \in \mathcal{A}$ as follows:

$$\theta_{ij} = r_{ij} - P_c d_{ij}.$$

Based on the above, the tramper problem in the base case reduces to that of finding the cycle that maximizes the profit per unit of time in the network $(\mathcal{N}, \mathcal{A})$ where crossing arc (i, j) earns θ_{ij} units of profit and requires τ_{ij} days. This maximum profit-to-time ratio problem has been analyzed in the literature; see Dantzig *et al.* (1966) for an early reference. When θ_{ij} are integers, Ahuja *et al.* (1993, Section 5.7) describe a binary search algorithm to find the optimal revenue per unit of time,

λ^D . The algorithm uses the fact that $\lambda^D \in [-\theta_{\max}, \theta_{\max}]$, where $\theta_{\max} = \max \{|\theta_{ij}| : (i, j) \in \mathcal{A}\}$. At a high level, the algorithm follows a sequence of iterations and after each such iteration, one is able to halve the interval of possible values for the optimal profit per unit of time. Each iteration starts with an interval of possible values $[\underline{\mu}, \bar{\mu}]$ and concludes whether $\lambda^D > (\underline{\mu} + \bar{\mu})/2$, $\lambda^D < (\underline{\mu} + \bar{\mu})/2$, or $\lambda^D = (\underline{\mu} + \bar{\mu})/2$. In the two first cases, one is able to restrict attention in the next iteration to an interval with half the length of the one one started with while in the last case, one has found the optimal profit per unit of time. At each iteration, one considers a network $(\mathcal{N}, \mathcal{A})$ with arc measures

$$\tilde{\theta}_{ij} = \theta_{ij} - \frac{\underline{\mu} + \bar{\mu}}{2} \tau_{ij},$$

and applies a shortest path label correcting algorithm to detect the existence or non-existence of a negative length cycle. If a negative length cycle exists, then $\lambda^D < (\underline{\mu} + \bar{\mu})/2$; if all cycles have positive length, then $\lambda^D > (\underline{\mu} + \bar{\mu})/2$; and if a zero length cycle exists, then $\lambda^D = (\underline{\mu} + \bar{\mu})/2$ and the cycle found is optimal in the original problem.

Overall, if $\tau_{\min} = \min \{\tau_{ij} : (i, j) \in \mathcal{A}\}$, it is possible to show that $O(\log(\tau_{\min} \theta_{\max}))$ iterations will suffice to find the optimal profit per unit of time and the associated cycle.

In the base case, refueling decisions are trivial and the optimal solution dictates to follow a cycle in the space of port locations. In particular, it is straightforward to establish that the cycles that maximize revenues and profits coincide. In that sense, the refueling and routing decisions are decoupled.

As highlighted in the introduction, an important feature associated with many networks of ports is that prices differ across ports. Next, we analyze how the structure of an optimal solution changes in the presence of price heterogeneity.

6.2.2 Deterministic Non-Uniform Fuel Prices

In this section, we assume again that prices are deterministic but we turn attention to the more general case in which prices differ across ports. Note that, while in the case of uniform prices the long-term optimal trajectory of a tramper is always a cycle on the graph $(\mathcal{N}, \mathcal{A})$, in the present case, the optimal trajectory need not be, in general, a cycle in the location space. However, as we show below, one can build on the analysis of the previous section to search for an optimal solution in the case of non-uniform prices, which can still be represented as a cycle, but in a generalized, location-inventory space. In the analysis below, it is convenient to introduce the following notation. Let $\tilde{\mathcal{N}} := \mathcal{N} \otimes \mathcal{I}$ denote the set of possible location-inventory combination pairs and let $\tilde{\mathcal{A}} := \{(z, z') : z \neq z', z, z' \in \tilde{\mathcal{N}}\}$ denote the set of all possible arcs connecting different elements of $\tilde{\mathcal{N}}$. A node z defines both a geographical location and an inventory position and hence moving from a node $z = (i, I)$ to a node $z' = (i', I')$ implies that the ship purchases $I' - I + d_{ii'}$ at

price P_i at node i . Note that $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ defines a complete graph built on $N \times (C+1)$ nodes. For any $z = (i, I) \in \tilde{\mathcal{N}}$, let $\tilde{O}(z) := \{z' = (i', I') : i' \in O(i), I - d_{ii'} \leq I' \leq C - d_{ii'}\}$ denote the set of nodes in $z' = (i', I') \in \tilde{\mathcal{N}}$ reachable in one transition from z . For each arc (z, z') belonging to $\tilde{\mathcal{A}}$, we can introduce the following measure

$$\tilde{\theta}_{z,z'} = r_{ii'} - P_i(I' - I + d_{ii'}), \quad (10)$$

that fully characterizes the profit contribution earned when moving from $z = (i, I)$ to $z' = (i', I')$. In addition, we let $\tilde{\tau}_{z,z'} = \tau_{ii'}$ for any $(z, z') \in \tilde{\mathcal{A}}$. Consider a cycle $\mathcal{C}$ on $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ consisting of the nodes $z_1 = (i_1, I_1)$, $z_2 = (i_2, I_2)$, ..., $z_k = (i_k, I_k)$, z_1 , where k is a positive integer, and let

$$\Theta_{\mathcal{C}} := \tilde{\theta}_{z_k, z_1} + \sum_{m=1}^{k-1} \tilde{\theta}_{z_m, z_{m+1}}, \quad (11)$$

$$T_{\mathcal{C}} := \tilde{\tau}_{z_k, z_1} + \sum_{m=1}^{k-1} \tilde{\tau}_{z_m, z_{m+1}}, \quad (12)$$

be the accumulated profit and the travel time associated with the cycle $\mathcal{C}$, respectively. Then the profit per unit of time associated with $\mathcal{C}$ is given by

$$\lambda_{\mathcal{C}} = \frac{\Theta_{\mathcal{C}}}{T_{\mathcal{C}}}. \quad (13)$$

The next result provides a generalization of the location connectivity assumption for $(\mathcal{N}, \mathcal{A})$ introduced in Section 3 and recasts the deterministic tramper problem as a problem of finding the cycle with maximum profit per unit of time on the network $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$.

Proposition 4 *a) Suppose that any two nodes $i \in \mathcal{N}$ and $j \in \mathcal{N} \setminus O(i)$ can be connected by a finite set of arcs in $\mathcal{A}$. Consider any cycle $\mathcal{C} : z_1 = (i_1, I_1)$, $z_2 = (i_2, I_2)$, ..., $z_k = (i_k, I_k)$, z_1 in $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ and any node $z \in \tilde{\mathcal{N}}$, $z \notin \mathcal{C}$. Then any point in the cycle is reachable from z in a finite number of transitions.*

b) For any point $z \in \tilde{\mathcal{N}}$ and any point z_i on a cycle $\mathcal{C}$, let $\mathcal{P}(z, z_i)$ denote a path from z to z_i with the smallest number of transitions. Further, let $F(z) = (i_F(z), I_F(z))$ denote the follower of point z on $\mathcal{P}(z, z_i)$, where $i_F(z) \in \mathcal{N}$ is the location of the next point and $I_F(z)$ is the inventory level corresponding to $F(z)$. Let $\mathcal{C}^ : z_1^* = (i_1^*, I_1^*)$, $z_2^* = (i_2^*, I_2^*)$, $z_k^* = (i_k^*, I_k^*)$, z_1^* be the cycle that achieves the highest profit per unit of time $\lambda_{\mathcal{C}^*}$ on the network $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$. Then, when at point $z = (i, I) \in \tilde{\mathcal{N}}$, the optimal policy for the tramper problem is given by*

$$j^*(z) = \begin{cases} i_{m+1}^*, & \text{if } z = z_m^*, m = 1, \dots, k-1, \\ i_1^*, & \text{if } z = z_k^*, \\ i_F(z), & \text{if } z \notin \{z_1^*, \dots, z_k^*\}, \end{cases} \quad (14)$$

$$q^*(z) = \begin{cases} I_{m+1}^* + d_{i, i_{m+1}^*} - I, & \text{if } z = z_m^*, m = 1, \dots, k-1, \\ I_1^* + d_{i, i_1^*} - I, & \text{if } z = z_k^*, \\ I_F(z) + d_{i, i_F(z)} - I, & \text{if } z \notin \{z_1^*, \dots, z_k^*\}, \end{cases} \quad (15)$$

and $\lambda = \lambda_{\mathcal{C}^*}$.

The proof of part a) of Proposition 4 is contained in the Appendix and provides an explicit method of reaching any point in the cycle z_i , $i = 1, \dots, k$ from any other point z . The result of part b) of Proposition 4 indicates that, under deterministic assumptions, a tramper should get to the “optimal” cycle $\mathcal{C}^*$ as fast as possible and then remain on it. Under such a policy, the profit values earned before the tramper reaches the cycle $\mathcal{C}^*$ do not contribute to the long-run value of the profit per unit of time - thus, the key element of the optimal location-refueling policy is the cycle $\mathcal{C}^*$.

In order to find the optimal cycle $\mathcal{C}^*$, the binary search algorithm discussed in the context of the base case can be adapted to the case of the expanded network $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$. The optimal solution is now a cycle in the generalized location-inventory space. The need to search for cycles beyond the location space is a consequence of the presence of different prices across ports.

6.2.3 Stochastic Uniform Prices

In this section, we analyze the impact of price stochasticity on the optimal solution structure. We assume that prices are uniform across all ports but that the common price P_c is stochastic. We illustrate through an example that the optimal solution need not be a cycle in the space of locations as in the base case.

Example 1: Consider a setting with three ports. Suppose that $r_{12} = r_{21} = r_{23} = r_{32} = \$20,000$ and $r_{13} = r_{31} = \$34,000$, $\tau_{12} = \tau_{21} = 1$, $\tau_{23} = \tau_{32} = 1$, $\tau_{13} = \tau_{31} = 2$, $d_{12} = d_{21} = 100$ mts, $d_{23} = d_{32} = 100$ mts, $d_{13} = d_{31} = 200$ mts, $C = 1,000$ mts. Suppose also that the price P_c can only take two values $P_L = \$50$ and $P_H = \$100$, and that the transition probability matrix is given by $[0.1, 0.9; 0.9, 0.1]$. Note that when P_c is equal to the time average of these two prices, $\$75$, the optimal routing reduces to the cycle $(1, 3)$. When P_c follows the two-value stochastic dynamics described above, the routing decisions, when in ports 1 and 3, become conditional on the state of the system (Table 2). In particular, when in port 3, the optimality of “go to port 1” decision now depends on either the observed price being low, or, if it is high, on the value of the on-board fuel inventory being also high enough to eliminate the need for refueling in the near future; in the cases when the fuel price is high and the need for refueling is strong, the best routing decision is to go to port 2. Since the problem data is symmetric, a similar routing policy applies when the vessel is in port 1.

As our example indicates, even in the absence of price dispersion across ports, a modest degree of stochasticity in fuel prices creates, in general, a substantial degree of complexity in optimal routing decisions, strongly coupling routing and refueling decisions.

Next port when in port 3											
On-Board Fuel Inventory (in 100's of mts)	0	1	2	3	4	5	6	7	8	9	10
Fuel Price = P_H	2	2	2	1	1	1	1	1	1	1	1
Fuel Price = P_L	1	1	1	1	1	1	1	1	1	1	1

Table 2: Optimal routing decision at port 3 as function of the fuel price and on-board fuel inventory

6.2.4 Coupling Between Routing and Refueling Decisions: The Role of Capacity

In this section we investigate the role of another factor, vessel capacity, on the coupling between routing and refueling decisions. We focus on the case of deterministic but non-uniform fuel prices. Let $\mathcal{C}^{(r)}$ be the cycle that maximizes revenues per unit of time (in the location space), where the revenues accumulated over a cycle are given by the sum of the rewards over that cycle. Let $\mathcal{C}^{(r)} : i_1, i_2, \dots, i_m, i_1^4$, $\bar{d} = \max_{(i,j) \in \mathcal{A}} d_{ij}$ and $i_{\min} = \arg \min_{i \in \mathcal{N}} \{P_i\}$. We suppose that $C > \bar{d}$. Consider the policy $\pi_{\mathcal{C}^{(r)}}$ that separates refueling and routing decisions as follows: it follows the cycle $\mathcal{C}^{(r)}$ until the reserve of fuel on board minus the fuel needed to reach the next port $I - d_{ii'}$ drops below $\bar{d}$. When this occurs, the vessel travels to port $i_{\min}$, replenishes up to C , and travels back to i and resumes following the route of the cycle. Let $m^* = \arg \min \{d_{i_k, i_{\min}}, k = 1, \dots, m\}$ and $I_{\mathcal{C}^{(r)}} = \sum_{i=1}^m d_{i, i+1} + d_{i_{m^*}, i_{\min}} + d_{i_{\min}, i_{m^*}}$. More formally, the policy can be defined as follows:

$$j^*(z) = \begin{cases} i_{l+1}, & \text{if } i = i_l \in \{i_1, i_2, \dots, i_m\} \setminus \{i_{m^*}\}, \\ i_{m^*+1}, & \text{if } i = i_{m^*} \text{ and } i_{\min} = i_{m^*}, \\ i_{\min}, & \text{if } i = i_{m^*}, :: i_{\min} \neq i_{m^*}, \text{ and } I < I_{\mathcal{C}^{(r)}} - d_{i_{\min}, i_{m^*}}, \\ i_{m^*+1}, & \text{if } i = i_{m^*}, :: i_{\min} \neq i_{m^*}, \text{ and } I \geq I_{\mathcal{C}^{(r)}} - d_{i_{\min}, i_{m^*}}, \\ i_{m^*}, & \text{if } i \notin \{i_1, \dots, i_m\}, \end{cases} \quad (16)$$

$$q^*(z) = \begin{cases} (d_{i_l, i_{l+1}} - I)^+, & \text{if } i = i_l \in \{i_1, i_2, \dots, i_m\} \setminus \{i_{m^*}\}, \\ (d_{i_{m^*}, i_{\min}} - I)^+, & \text{if } i = i_{m^*}, :: i_{\min} \neq i_{m^*}, \text{ and } I < I_{\mathcal{C}^{(r)}} - d_{i_{\min}, i_{m^*}}, \\ 0, & \text{if } i = i_{m^*}, :: i_{\min} \neq i_{m^*}, \text{ and } I \geq I_{\mathcal{C}^{(r)}} - d_{i_{\min}, i_{m^*}}, \\ (d_{i, i_{m^*}} - I)^+, & \text{if } i \notin \{i_1, \dots, i_m\} \cup \{i_{\min}\}, \\ C - I, & \text{if } i = i_{\min}, \end{cases} \quad (17)$$

The next result establishes the near-optimality of this policy in cases where the capacity is large relative to the maximum fuel consumption required for a single inter-port transition. In particular, define

$$\zeta = \frac{\max_{(i,j) \in \mathcal{A}} d_{ij}}{C}$$

as the proportion of the total capacity needed for the most fuel-consuming inter-port transition. We have the following result.

⁴We adopt the convention that i_{m+1} refers to i_1 to simplify notation.

Proposition 5 *Let $\lambda^D > 0$ be the optimal profit per unit of time. Suppose that $C \geq I_{\mathcal{C}(r)}$. Then, the long run profit per unit of time associated with the policy $\pi_{\mathcal{C}(r)}$, $\lambda_{\pi_{\mathcal{C}(r)}}$, satisfies:*

$$\frac{\lambda_{\pi_{\mathcal{C}(r)}}}{\lambda^D} \geq \frac{1 - 4\zeta}{1 - 2\zeta}. \quad (18)$$

Hence, as ζ converges to zero, one can essentially decouple refueling and routing decisions. Having ζ “small” corresponds to a tramper operating in a region with a high number of ports that are close to each other. This would be the case for a tramper restricting its operation to a small region of the world, such as e.g., the Mediterranean. For $\zeta = 1/8$, one can lower bound the ratio of the profit of the proposed policy to the optimal profit by $2/3$ and for $\zeta = 1/12$, by $4/5$.

6.3 Heuristic Profit Management Policies and Their Performance

In practice, even small-size instances of the dynamic program (8)-(9) are often computationally intractable. In particular, it is nearly impossible to solve such a problem to optimality for a network with more than 3 ports since the size of the state space grows exponentially with the number of ports. In such a setting, developing heuristics with lower computational requirements becomes critical.

In addition, as illustrated in the previous section, the resulting optimal profit management policies may take a complex and nonintuitive form in the general case. This section focuses on the development of insights on the optimal policy and of two potential heuristic profit management policies.

6.3.1 The Optimal Location Cycle and the Optimal Location-Inventory Cycle Heuristics

The first profit management heuristic policy uses the approach described in section 6.2.1 to decouple the routing and refueling decisions. For any port $i \in \mathcal{N}$, let $\underline{P}_i$ denote the smallest value that the price at port i can take and let $\underline{P} = \min_{i \in \mathcal{N}} \underline{P}_i$. The heuristic proceeds in two steps. In the first step, it is assumed that all fuel prices are constant over time and equal to a common value $P_c = \underline{P}$ and all arc rewards are constant equal to their expectation; using this simplification, the profit maximizing location cycle is determined. In the second step, the optimal refueling policy is determined for a liner following that location cycle. We refer to this heuristic as the “OLC” for Optimal Location Cycle. In summary, the OLC heuristic tackles the combinatorial nature of the dynamic routing decisions by separating them from the refueling decisions. It is important to note that the OLC heuristic still requires to solve a liner problem over a given cycle, which might imply significant computational requirements if the cycle has a large number of ports.

Below, we provide a bound for the performance of the OLC heuristic. For any arc $(i, j) \in \mathcal{A}$, let $\bar{r}_{ij}$ denote the largest possible value of the reward on arc (i, j) . In addition, let $\mathcal{S}_C$ denote the set of all cycles in the location space $(\mathcal{N}, \mathcal{A})$ and define

$$\begin{aligned}\Delta_r &= \max_{(i,j) \in \mathcal{A}} \{\bar{r}_{ij} - \mathbb{E}[r_{ij}]\}, \\ \Delta_p &= \max_{i \in \mathcal{N}} \{\mathbb{E}_\infty[P_i] - \underline{P}\}, \\ \omega &= \max_{\mathcal{C} \in \mathcal{S}_C} \frac{\sum_{(i,j) \in \mathcal{C}} 1}{\sum_{(i,j) \in \mathcal{C}} \tau_{ij}}, \\ \rho &= \max_{\mathcal{C} \in \mathcal{S}_C} \frac{\sum_{(i,j) \in \mathcal{C}} d_{ij}}{\sum_{(i,j) \in \mathcal{C}} \tau_{ij}},\end{aligned}$$

where $\mathbb{E}_\infty$ corresponds to the expectation with respect to the steady-state distribution of prices.

Proposition 6 *Let λ^* and λ_{OLC} denote the profit per unit of time values achieved by an optimal policy and by the OLC heuristic, respectively. Then*

$$\lambda^* - \lambda_{OLC} \leq \omega \Delta_r + \rho \Delta_p. \quad (19)$$

The bound presented in Proposition 6 is sum of two components. Each of the components is the product of two terms where the first one, ω (respectively ρ) depends exclusively on the characteristics of the network of ports. Δ_r represents the influence of the stochastic nature of rewards while Δ_p represents the influence of the price heterogeneity across ports. Note that the bound (19) emphasizes the near-optimality of the OLC heuristic in settings where rewards are mildly stochastic and where prices do not significantly differ across ports, i.e., where Δ_r and Δ_p are small compared to the average values of rewards and prices, respectively. It is also important to note that the decoupling between routing and refueling followed by this heuristic implies that when there is significant price heterogeneity across ports (even if prices are constant across time), ones runs the risk of “missing” the ports where fuel is sold at particularly low prices, which might have severe impacts on performance.

The second heuristic policy we consider is based on the analysis conducted for the deterministic case in section 6.2.2. In particular, consider the maximum profit-to-time ratio cycle in the network $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ with profits

$$\tilde{\theta}_{z,z'} = \mathbb{E}[r_{ii'}] - \mathbb{E}_\infty[P_i](I' - I + d_{ii'}). \quad (20)$$

Let π_{OLIC} denote the policy which leads the vessel through this location-inventory cycle and let $\mathcal{P}_{OLIC}$ denote the subset of stationary policies that form cycles in the location-inventory network. The analysis performed in section 6.2.2 implies that for the general stochastic tramper problem, π_{OLIC} achieves the highest possible long-term average profit among all policies in $\mathcal{P}_{OLIC}$. In summary, the optimal location-inventory cycle policy (OLIC) captures both the heterogeneity in prices

as well as the combination of route selection with refueling decisions but ignores stochastic variations of prices and rewards. Note that this feature of the OLIC policy is in contrast with the OLC heuristic that captures price stochasticity across the ports it visits.

Below, we provide a bound for the performance of the OLIC heuristic. In addition to the notation introduced above, let $\tilde{\mathcal{S}}_C$ denote the set of all cycles in the location-inventory space $(\tilde{\mathcal{N}}, \tilde{\mathcal{A}})$ and define

$$\begin{aligned}\tilde{\Delta}_p &= \max_{i \in \mathcal{N}} \{\mathbb{E}_\infty[P_i] - \underline{P}_i\}, \\ \tilde{\omega} &= \max_{\mathcal{C} \in \tilde{\mathcal{S}}_C} \frac{\sum_{(z, z') \in \mathcal{C}} 1}{\sum_{(z, z') \in \mathcal{C}} \tau_{ii'}}, \\ \tilde{\rho} &= \max_{\mathcal{C} \in \tilde{\mathcal{S}}_C} \frac{\sum_{(z, z') \in \mathcal{C}} d_{ii'}}{\sum_{(z, z') \in \mathcal{C}} \tau_{ii'}}.\end{aligned}$$

Proposition 7 *Let λ^* and λ_{OLIC} denote the profit per unit of time values achieved by an optimal policy and by the OLIC heuristic, respectively. Then*

$$\lambda^* - \lambda_{OLIC} \leq \tilde{\omega} \Delta_r + \tilde{\rho} \tilde{\Delta}_p. \quad (21)$$

Note that the bound (21) emphasizes the near-optimality of the OLIC heuristic in mildly stochastic settings, i.e., where Δ_r and $\tilde{\Delta}_p$ are small compared to the average values of rewards and prices, respectively. Contrasting the OLIC and OLC heuristics, one observes that the former captures price heterogeneity across ports and is not exposed to the risk of potentially ignoring ports where fuel is sold at a discount; this is reflected in the performance bounds through the dependence on $\tilde{\Delta}_p$ for the OLIC heuristic (which is driven exclusively by price stochasticity) as opposed to Δ_p for OLC heuristic (which is driven by price heterogeneity).

While we will restrict attention to the two heuristics presented above, it is worth noting that the analysis of section 6.2.4 can be a basis for the development of additional heuristics, where one could determine a cycle that maximizes average accumulated rewards per unit of time and then only deviate from such a cycle to replenish at ports where fuel prices are heavily discounted.

6.3.2 Comparing the Performance of the Heuristics

In this section we report the results of a numerical study designed to gain insights on the optimal policy and the performance of the two heuristics. The problem setting for our numerical study is designed to test heuristics in a real-life setting representing a tramper traversing a small network of 3 ports. The choice of a small size network for testing purposes is dictated by the computational challenges associated with obtaining the optimal policies for problems with more than three ports (the heuristics can be applied to real-life problems with a larger number of ports). For example, a

simple problem instance with four ports, 15 crude oil price levels, 10 possible fuel inventory levels, 2 levels of local Markovian price corrections at each port, 2 levels of local iid price corrections, and 2 revenue levels on each arc has 38,400 states. Such a problem instance takes several days to solve using our tool of choice, Mathematica (it takes around 10 hours to solve a typical tramper problem with 3 ports.)

We consider a tramper with capacity of 6000 mts operating between the ports of Houston (port 1), Rotterdam (port 2) and Abidjan (port 3). Figures 2 and 3 show a schematic representation of the distances between these ports expressed in nautical miles and of the corresponding network parameters, respectively. We note that the distances between the ports are calculated assuming that the tramper follows the standard inter-port trajectories passing through the so-called “junction points” as specified by the National Imagery and Mapping Agency (2001). For example, in the case of a travel between Rotterdam and Houston, a tramper vessel would pass from Rotterdam to the junction point at Ile de Ouessant, then to the junction point at Great South Banks, followed by the junction point at the Straits of Florida, and, finally, by the port of Houston, resulting in a travel distance of 5009 nautical miles. The travel times and fuel consumption values for the trips between ports were calculated using the typical average travel speed of 15 nautical miles per hour and the typical average fuel consumption rate of 75 mts per day. The resulting travel time values were rounded to the nearest day and the resulting fuel consumption values - to the nearest multiple of 200 mts. Note that we assume, for the sake of simplicity, that the travel times and the fuel consumption values are symmetric, i.e., $\tau_{ij} = \tau_{ji}$ and $d_{ij} = d_{ji}$ for all $i, j = 1, 2, 3, i \neq j$. In summary, the fuel consumptions and travel times between the 3 ports are given by

$$\mathbf{d} = \begin{bmatrix} 0 & 1000 & 1200 \\ 1000 & 0 & 800 \\ 1200 & 800 & 0 \end{bmatrix}, \quad \boldsymbol{\tau} = \begin{bmatrix} 0 & 14 & 15 \\ 14 & 0 & 11 \\ 15 & 11 & 0 \end{bmatrix}. \quad (22)$$

To reduce the computational effort, we assume that the ship can only replenish fuel in multiples of 200 mts. We use the price dynamics model outlined in (3) and the parameter estimation procedure described in Appendix B. As before, the value of the parameter $\hat{\gamma}$ representing a conversion factor between the time average of the oil price expressed in \$/bbl and the time-and-location average of bunkers prices expressed in \$/mts is estimated as 4.43, and the values of other parameter estimates are presented in Table 3.

We observe that, as Table 3 indicates, the port of Rotterdam has, on average, a pronounced price advantage over the other two ports, while refueling at Abidjan is, in general, undesirable. As in the liner case, the dynamics of oil price P_0 are described by (B9) in Appendix B. This parameter combination results in Markovian price dynamics under which the high degree of inertia in price values and high variance in supply corrections allow for various orderings of prices among the three ports to be realized. We use a simple form for random price correction terms by assuming that ε_i

Figure 2: Schematic geographical representation of the distances between ports (in nautical miles).

Figure 3: Travel times and fuel consumption values for ports.

Parameter	Houston (Port 1)	Rotterdam (Port 2)	Abidjan (Port 3)
A , \$/mts	-13.42	-28.44	9.62
δ , \$/mts	31.71	29.72	19.21
η	0.93	0.99	0.90
σ , \$/mts	13.94	14.83	11.91

Table 3: Parameters of the fuel price dynamics for the ports of Houston, Rotterdam and Abidjan.

at port i takes two values, $\pm\sigma_i$, w.p. 0.5 for each value. The oil price dynamics is taken to follow the mean-reverting model (B9) with $\sigma_0 = 1.25$, $\lambda = 0.96$, $P_0^{\min} = 46$ and $P_0^{\max} = 62$.

In practice, the revenues associated with transporting cargo between ports are established in the process of negotiations, their exact values being closely guarded commercial information. In our numerical study, we assume a simple setting with constant revenue matrix

$$\mathbf{r} = \begin{bmatrix} 0 & r_{12} & r_{13} \\ r_{21} & 0 & r_{23} \\ r_{31} & r_{32} & 0 \end{bmatrix}, \quad (23)$$

where r_{ij} , $i \neq j$ are taken to be proportional to the distances between ports, $r_{ij} = \kappa_{ij}d_{ij}$, with κ_{ij} being the parameters that we varied. We take $\kappa_{12} = \kappa_{21} = \kappa_{13} = \kappa_{31} = \kappa_{23} = \kappa_{32} = 250$ as the base scenario for our study. Note that for such values of revenue parameters the profit rate generated by the optimal tramper policy turns out to be \$794/day, a very modest amount roughly corresponding to the price of 4 metric tons of fuel, or less than 5% of typical daily fuel cost. Thus, the base scenario value of revenue parameters corresponds to a reasonable lower bound on revenues which guarantee profitability.

In our numerical study we have used a standard value iteration approach to solve the tramper dynamic program. Table 4 illustrates the optimal route selection decisions for the base scenario in the states with $P_0 = \$54$ and zero on-board fuel inventory. In this table we have used a short-hand notation “ $\pm\delta_i \pm \sigma_i$ ” to denote the price state $P_i = \gamma P_0 + A_i \pm \delta_i \pm \sigma_i$ at the “current” port $i = 1, 2, 3$, (i.e., at the port where the routing decision is made) and “ $\pm\delta_i$ ” to denote the price state at other ports.

As Table 4 shows, the resolution of a trade-off between low fuel prices and high revenue values often takes an intuitive form. For example, when in Houston, the choice between higher revenue rate associated with the Houston-Abidjan link and cheaper fuel in Rotterdam is resolved as follows: when the fuel price in Houston is at its lowest possible level, immediate refueling allows the tramper to benefit from the revenue advantage of going to Abidjan; on the other hand, when the fuel price in Houston is at its highest possible level, the prospective of cheaper refueling in Rotterdam outweighs the revenue considerations. For the intermediate values of fuel prices at Houston, the trade-off

Next Port at Houston	Price at Houston			
Prices at Rotterdam and Abidjan	$-\delta_1 - \sigma_1$	$+\delta_1 - \sigma_1$	$-\delta_1 + \sigma_1$	$+\delta_1 + \sigma_1$
$(-\delta_2, -\delta_3)$	Abidjan	Rotterdam	Rotterdam	Rotterdam
$(-\delta_2, +\delta_3)$	Abidjan	Rotterdam	Rotterdam	Rotterdam
$(+\delta_2, -\delta_3)$	Abidjan	Rotterdam	Abidjan	Rotterdam
$(+\delta_2, +\delta_3)$	Abidjan	Rotterdam	Abidjan	Rotterdam

Next Port at Rotterdam	Price at Rotterdam			
Prices at Houston and Abidjan	$-\delta_2 - \sigma_2$	$+\delta_2 - \sigma_2$	$-\delta_2 + \sigma_2$	$+\delta_2 + \sigma_2$
$(-\delta_1, -\delta_3)$	Abidjan	Houston	Abidjan	Houston
$(-\delta_1, +\delta_3)$	Abidjan	Houston	Abidjan	Houston
$(+\delta_1, -\delta_3)$	Abidjan	Houston	Abidjan	Houston
$(+\delta_1, +\delta_3)$	Abidjan	Houston	Abidjan	Houston

Next Port at Abidjan	Price at Abidjan			
Prices at Houston and Rotterdam	$-\delta_3 - \sigma_3$	$+\delta_3 - \sigma_3$	$-\delta_3 + \sigma_3$	$+\delta_3 + \sigma_3$
$(-\delta_1, -\delta_2)$	Rotterdam	Rotterdam	Rotterdam	Rotterdam
$(-\delta_1, +\delta_2)$	Houston	Houston	Houston	Houston
$(+\delta_1, -\delta_2)$	Rotterdam	Rotterdam	Rotterdam	Rotterdam
$(+\delta_1, +\delta_2)$	Houston	Houston	Houston	Rotterdam

Table 4: Optimal route selection decisions for the base scenario in the states with $P_0 = \$54$ and zero on-board fuel inventory.

becomes more subtle: Rotterdam is preferred unless the price in Houston is cheap enough to make immediate refueling in Houston more advantageous than the later refueling in Rotterdam, and, thus, to allow the routing decision to be based solely on revenue rate consideration. The subtlety of such a trade-off is clearly illustrated in the optimal routing decisions when in Abidjan: while in the majority of cases the next-port-to-visit is decided based on the basis of the expected fuel price, in the setting where the prices at both Houston and Rotterdam are expected to remain high, the situation is “too close to call” without accounting for the fuel price in Abidjan itself. In particular, as the last line of the Table 4 indicates, the slight expected price advantage of Rotterdam is canceled out by a slight revenue rate advantage associated with Abidjan-Houston link, unless the price in Abidjan is at its highest level as well.

In our numerical study we consider four different variations of the base scenario, each variation designed to emphasize the influence of a particular location cycle within the 3-port network we consider. In particular, in the first variation of the base scenario we increase the revenue values associated with Houston-Rotterdam-Houston routing cycle by considering $\kappa_{12} = \kappa_{21} =$

$\kappa_{12} = \kappa_{21}$	ε_{OLC}	Locations visited: OLC	$\varepsilon_{\text{OLIC}}$	Locations visited: OLIC
250	3.1	(1, 2)	51.5	(1, 3, 2)
275	0.0	(1, 2)	37.8	(1, 2)
300	0.0	(1, 2)	28.7	(1, 2)
325	0.0	(1, 2)	23.1	(1, 2)
350	0.0	(1, 2)	19.4	(1, 2)
$\kappa_{13} = \kappa_{31}$	ε_{OLC}	Locations visited: OLC	$\varepsilon_{\text{OLIC}}$	Locations visited: OLIC
250	3.1	(1, 2)	51.5	(1, 3, 2)
275	13.1	(1, 3)	24.1	(1, 3)
300	3.6	(1, 3)	28.7	(1, 3)
325	0.0	(1, 3)	12.1	(1, 3)
350	0.0	(1, 3)	6.8	(1, 3)
$\kappa_{23} = \kappa_{32}$	ε_{OLC}	Locations visited: OLC	$\varepsilon_{\text{OLIC}}$	Locations visited: OLIC
250	3.1	(1, 2)	51.5	(1, 3, 2)
275	0.0	(2, 3)	29.7	(2, 3)
300	0.0	(2, 3)	22.1	(2, 3)
325	0.0	(2, 3)	17.7	(2, 3)
350	0.0	(2, 3)	14.7	(2, 3)
$\kappa_{12} = \kappa_{23} = \kappa_{31}$	ε_{OLC}	Locations visited: OLC	$\varepsilon_{\text{OLIC}}$	Locations visited: OLIC
250	3.1	(1, 2)	51.5	(1, 3, 2)
275	0.9	(1, 2, 3)	37.2	(1, 2, 3)
300	0.4	(1, 2, 3)	28.9	(1, 2, 3)
325	0.2	(1, 2, 3)	24.0	(1, 2, 3)
350	0.1	(1, 2, 3)	20.7	(1, 2, 3)

Table 5: Relative performance gaps between the optimal policy and the “optimal location cycle” ($\varepsilon_{\text{OLC}} = (\lambda_{\text{opt}} - \lambda_{\text{OLC}})/\lambda_{\text{opt}} \times 100\%$) and the “optimal location-inventory cycle” ($\varepsilon_{\text{OLIC}} = (\lambda_{\text{opt}} - \lambda_{\text{OLIC}})/\lambda_{\text{opt}} \times 100\%$) heuristics.

250, 275, ..., 350 (the rest of revenue values remain at their base scenario values). Similarly, the second variation raises the importance of Houston-Abidjan-Houston routing cycle by setting $\kappa_{13} = \kappa_{31} = 250, 275, \dots, 350$. The third and forth settings emphasize Rotterdam-Abidjan-Rotterdam and Houston-Rotterdam-Abidjan-Houston routing cycle by setting $\kappa_{23} = \kappa_{32} = 250, 275, \dots, 350$ and $\kappa_{12} = \kappa_{23} = \kappa_{31} = 250, 275, \dots, 350$, respectively. Note that in the latter case, the rewards on arcs are not anymore symmetric. The relative performance of the two heuristic profit management approaches defined above (optimal location cycle and optimal location-inventory cycle) under

these four problem variations is illustrated in Table 5. As we can see, both heuristics “latch on” the appropriate route in the majority of cases we studied. The optimal location cycle heuristic exhibits a very robust performance across all the problem instances - with the worst-case relative performance gap of 13.1%. However, the best location-inventory cycle heuristic falls far behind in its performance even in cases when the location component of the location-inventory cycle is properly identified. Such inadequate performance is to be expected in settings where the stochasticity of the fuel prices is pronounced and the failure to capitalize on the optimal re-fueling policies significantly hinders the profit generation process. Indeed, the OLC heuristic adopts an optimal (state-dependent) refueling policy over the cycle that it selected while the OLIC heuristic refuels according to what is prescribed by the location-inventory cycle, independently of current price realizations. At the same time, it is also important to note that the cases with particularly high values of the relative performance gap stem from the low-profitability settings (i.e., where λ_{opt} is small) in which we test the heuristics.

7 Summary and Future Work

In this study we have developed a profit optimization model for a marine shipping company which owns a fleet of liners (vessels with fixed route) and trampers (vessels for which the route can be selected in a dynamic fashion). A model was also proposed for the dynamics of bunkers prices and was calibrated using the actual pricing data for a network of 18 ports for a period of 6 months. For the liner case, we formulate the vessel refueling problem as a long-term average stochastic dynamic program and prove that the optimal refueling policy has a capacitated price-dependent buy-up-to form. The monotonicity of the optimal fuel-up-to levels is established under additional assumptions on the stochastic properties of the fuel price dynamics. In the tramper case, the task of refueling is blended with a combinatorial route selection. We develop insights based on several special cases that allow one to isolate the impact of the features associated with this problem: heterogeneity of prices across ports, stochasticity of prices, and capacity constraints.

The single-vessel approach developed in this paper can serve as a good initial step to optimizing fleet profits. At the same time, real business settings are often characterized by the presence of volume refueling discounts which can only be fully exploited using more than one vessel. Thus, the development of multiple-vessel profit management models represents a challenging research direction which will be of immediate interest to practitioners.

Acknowledgments

The authors thank three anonymous reviewers and an associate editor for providing many valuable suggestions that greatly improved the paper.

References

- Ahuja R.K., T.L. Magnanti, and J.B. Orlin. 1993. “Network Flows”, Prentice Hall.
- Asche, F., O. Gjølborg, and T. Völker. 2003. “Price Relationships in the Petroleum Market: an Analysis of Crude Oil and Refined Product Prices”, *Energy Economics* **25**, 289-301.
- Bertsekas, D.P. 2000. “Dynamic Programming and Optimal Control“, Volume II, Athena Scientific.
- Borenstein, S., A.C. Cameron, and R. Gilbert. 1997. “Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?”, *The Quarterly Journal of Economics*, **112** (1), 305-339.
- Brennan, M.J., and E.S. Schwartz. 1985. “Evaluating Natural Resource Investments”, *Journal of Business* **58**, 135-157.
- Christiansen, M., K. Fagerholt, and D. Ronen. 2004. “Ship Routing and Scheduling: Status and Perspectives”, *Transportation Science* **38**, 1-18.
- Christiansen, M., K. Fagerholt, B. Nygreen, and D. Ronen. 2005. “Maritime Transportation”, a chapter in *Handbook for Operations Research and Management Science: Transportation*, edited by Barnhart and Laporte, North-Holland, Amsterdam.
- Cho, S.-C., A.N. Perakis. 1996. “Optimal Liner Fleet Routing Strategies”, *Maritime Policy and Management* **23**, 249-259.
- Cortazar, G., and E.S. Schwartz. 1994. “The Evaluation of Commodity Contingent Claims”, *Journal of Derivatives* **1**, 27-39.
- Dantzig, G.B., W. Blattner, and M.R. Rao. 1966. “Finding a Cycle in a Graph with Minimum Cost to Time Ratio with Applications to a Ship Routing Problem”, *Theory of Graphs International Symposium, Dunod, Paris, and Gordon and Breach, New York*, 77-83.
- Dantzig, G.B., and D.R. Fulkerson. 1954. “Minimizing a Number of Tankers to Meet a Fixed Schedule”, *Naval Res. Logist. Quart.* **1**, 217-222.
- Darnell, D.W., and C. Loffin. 1977. “National Airlines Fuel Management and Allocation Model”, *Interfaces* **7**, 1-16.
- Dixit, A.K., and R.S. Pindyck. 1994. *Investment Under Uncertainty*, Princeton University Press, Princeton, NJ.
- Either, R., and T. Mount. 1998. “Estimating Volatility of Spot Prices in Restructured Electricity Markets and the Implications for Option Values”, working paper, Cornell University.

- Gendreau, M., Laporte, G. and R. Deguin. 1995. "Stochastic Vehicle Routing", *European Journal of Operational Research* **88**, 3-12.
- Hamilton, J.D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica* **57**, 357-384.
- Hamilton, J.D., and R. Susmel. 1994. "Autoregressive Conditional Heteroscedasticity and Changes in Regime", *Journal of Econometrics* **64**, 307-333.
- Imai, A., F. Rivera. 2001. "Strategic Fleet Size Planning for Maritime Refrigerated Containers", *Maritime Policy and Management* **28**, 361-374.
- Kalymon, B.A. 1971. "Stochastic Prices in a Single-Item Inventory Purchasing Model", *Operations Research* **19**, 1434-1458.
- Kaufmann, R.K., and C. Laskowski. 2005. "Causes for an Asymmetric Relation between The Price of Crude Oil and Refined Petroleum Products", *Energy Policy* **33**, 1587-1596.
- Laughton, D.G., and H.D. Jacoby. 1995. "The Effects of Reversion on Commodity Projects of Different Length", in *Real Options in Capital Investment: Models, Strategies, and Applications*, edited by L. Trigeorgis. Praeger, Westport, CT, 185-205.
- Marine Transportation System National Advisory Council. U.S. Economic Growth and the Marine Transportation System. Washington, DC: U.S. Department of Transportation, December 18, 2000.
- Moinzadeh, K. 1997. "Replenishment and Stocking Policies for Inventory Systems with Random Deal Offerings", *Management Science* **43**, 334-342.
- National Imagery and Mapping Agency. 2001. Publication 151: Distances Between Ports. Bethesda, MD (www.nga.mil, last checked on 01/04/2008).
- Noel, Michael D. 2007a. "Edgeworth Price Cycles: Evidence from the Toronto Retail Gasoline Market", *Journal of Industrial Economics* **55**, 69-92.
- Noel, Michael D. 2007b. "Edgeworth Price Cycles, Cost-based Pricing and Sticky Pricing in Retail Gasoline Markets", *Review of Economics and Statistics* **89**, 324-334.
- Ozekici, S., and M. Parlar. 1999. "Inventory Models with Unreliable Suppliers in a Random Environment", *Annals of Operations Research* **91**, 123-136.
- Paddock, J.L., D.R. Siegel, and J.L. Smith. 1988. "Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases", *Quarterly Journal of Economics* **103**, 479-503.
- Perakis, A. N., and D.I. Jaramillo. 1991. "Fleet Deployment Optimization for Liner Shipping, Part 1: Background, Problem Formulation and Solution Approaches", *Maritime Policy and Management* **18**, 183-200.
- Powell, B. J., and A.N. Perakis. 1997. "Fleet deployment Optimization for Liner Shipping: An Integer Programming Model", *Maritime Policy and Management* **24**, 183-192.

- Richetta, O., R.C. Larson. 1997. “Modeling the Increased Complexity of New York City’s Refuse Marine Transport System”, *Transportation Science* **31**, 272–293.
- Scarf, H. 1960. “The Optimality of (S, s) Policies in the Dynamic Inventory Problem”, *Mathematical Methods in Social Sciences*, Chapter 13, edited by Arrow, Karlin, and Suppes, Stanford University Press, Stanford, CA.
- Smith, J.E., and K.F. McCardle. 1998. “Valuing Oil Properties: Integrating Option Pricing and Decision Analysis Approaches”, *Operations Research* **46**, 198-217.
- Smith, J.E., and K.F. McCardle. 1999. “Options in the Real World: Lessons Learned in Evaluating Oil and Gas Investments”, *Operations Research* **47**, 1-15.
- Song, J.-S., and P. Zipkin. 1993. “Inventory Control in a Fluctuating Demand Environment”, *Operations Research* **41**, 351-370.
- Stroup, J.S., and R.D. Wollmer. 1992. “A Fuel Management Model for the Airline Industry”, *Operations Research* **40**, 229-237.
- US Census Bureau, 2002 Economic Census, Transportation and Warehousing Industry Series. Washington, DC: US Department of Commerce.

Appendices for “Going Bunkers: The Joint Route Selection and Refueling Problem”

A Proofs

Proof of Proposition 1

We establish the result using value iteration. We introduce a modified system to obtain an equivalent discrete-time average cost problem (see Bertsekas(2000)). We let β denote a constant in $(0, \min_i \tau_{ij(i)})$, $\tilde{w}(i, I, \mathbf{P}) = w(i, I, \mathbf{P}) / \beta$ and $\tilde{P}_i = P_i / \tau_{ij(i)}$. In addition, we let $\nu_{(i, I, \mathbf{P}), (j, I', \mathbf{P}')} (u)$ denote the transition probability from a state $(i, I, \mathbf{P})$ to a state $(j, I', \mathbf{P}')$ under a given replenishment decision u . We now define $\tilde{\nu}_{(i, I, \mathbf{P}), (j, I', \mathbf{P}')} (u) = (\beta / \tau_{ij(i)}) \nu_{(i, I, \mathbf{P}), (j, I', \mathbf{P}')} (u)$ for $(i, I, \mathbf{P}) \neq (j, I', \mathbf{P}')$ and $\tilde{\nu}_{(i, I, \mathbf{P}), (i, I, \mathbf{P})} (u) = 1 - (\beta / \tau_{ij(i)})$, and let $\tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}}$ denote the expectation under $\tilde{\nu}$. Then μ and $\tilde{w}$ satisfy

$$\tilde{w}(i, I, \mathbf{P}) = -\mu + \min_{q \in Q(i, j, I)} \left(\tilde{P}_i q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}(j(i), I + q - d_{ij(i)}, \mathbf{P}') \right] \right),$$

and the optimal policies are identical for both problems. Now, relative value iteration is guaranteed to converge to the solution of the Bellman equation since the modified system is aperiodic (cf. Prop 4.3.1 in Bertsekas(2000)). We let $\tilde{w}_k$ denote the sequence of functions obtained through relative value iteration. The recursion takes the following form:

$$\begin{aligned} \tilde{w}_{k+1}(i, I, \mathbf{P}) &= \min_{q \in Q(i, j, I)} \left(\tilde{P}_i q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k(j(i), I + q - d_{ij(i)}, \mathbf{P}') \right] \right) \\ &\quad - \min_{q \in Q(i_0, j(i_0), I_0)} \left(\tilde{P}_{i_0} q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}_0} \left[\tilde{w}_k(j(i_0), I_0 + q - d_{i_0 j(i_0)}, \mathbf{P}') \right] \right), \end{aligned} \quad (\text{A1})$$

where (i_0, I_0) is an arbitrary location/inventory state. We have

$$\lim_{k \rightarrow \infty} \tilde{w}_k(i, I, \mathbf{P}) = \tilde{w}(i, I, \mathbf{P}), \quad (\text{A2})$$

$$\mu = \lim_{k \rightarrow \infty} \min_{q \in Q(i_0, j(i_0), I_0)} \left(\tilde{P}_{i_0} q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}_0} \left[\tilde{w}_k(j(i_0), I_0 + q - d_{i_0 j(i_0)}, \mathbf{P}') \right] \right). \quad (\text{A3})$$

We next show that if $\tilde{w}_k(i, I, \mathbf{P})$ is a piecewise linear convex function with respect to I with changes of slope that belong to $B(i)$, then so is $\tilde{w}_{k+1}(i, I, \mathbf{P})$.

Suppose that $\tilde{w}_k(i, I, \mathbf{P})$ is a piecewise linear convex function with respect to I with changes of slope that belong to $B(i)$. Consider equation (A1) defining $\tilde{w}_{k+1}(i, I, \mathbf{P})$, the second term is independent of I and hence we need only focus on the first term. Define

$$f_k(i, y, \mathbf{P}) := \tilde{P}_i y + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k(j(i), y - d_{ij(i)}, \mathbf{P}') \right], \quad (\text{A4})$$

$$\begin{aligned} g(i, I, \mathbf{P}) &:= \min_{q \in Q(i, j, I)} \left(\tilde{P}_i q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k(j(i), I + q - d_{ij(i)}, \mathbf{P}') \right] \right) \\ &= -\tilde{P}_i I + \min_{d_{ij(i)} \leq y \leq C} (f_k(i, y, \mathbf{P})). \end{aligned} \quad (\text{A5})$$

Since $\tilde{w}_k(j(i), y, \mathbf{P}')$ is piecewise linear convex in y with changes of slopes in $B(j(i))$, $f_k(i, y, \mathbf{P})$ is piecewise linear convex (with respect to y) with changes of slope in the set $\{b + d_{ij(i)} : b \in B(j(i)), b + d_{ij(i)} \leq C\} = B(i)$. Let $y_k^*(i, P)$ denote a minimizer in (A5); it is possible to select it such that it occurs at a slope change, i.e., $y_k^*(i, P) \in B(i)$. We then have

$$\begin{aligned}
& g(i, I, \mathbf{P}) \\
&= \tilde{P}_i(y_k^*(i, P) - I)^+ + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k \left(j(i), I + (y_k^*(i, P) - I)^+ - d_{ij(i)}, \mathbf{P}' \right) \right] \\
&= 1\{I \leq y_k^*(i, P)\} \left[\tilde{P}_i(y_k^*(i, P) - I) + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k \left(j(i), y_k^*(i, P) - d_{ij(i)}, \mathbf{P}' \right) \right] \right] \\
&\quad + 1\{I > y_k^*(i, P)\} \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k \left(j(i), I - d_{ij(i)}, \mathbf{P}' \right) \right] \tag{A6}
\end{aligned}$$

In the expression above, the first term is clearly piecewise linear with changes of slope contained in $B(i)$ since $y_k^*(i, P) \in B(i)$. Now the same holds for the second term by the assumption on $\tilde{w}_k$. We deduce that $g(i, \cdot, \mathbf{P})$ is piecewise linear with changes of slope contained in $B(i)$. Next, we check that $g(i, \cdot, \mathbf{P})$ is convex.

$$\begin{aligned}
& g(i, I + 1, \mathbf{P}) - g(i, I, \mathbf{P}) \\
&= -1\{I + 1 \leq y_k^*(i, P)\} \tilde{P}_i \\
&\quad + 1\{I + 1 > y_k^*(i, P)\} \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k \left(j(i), I + 1 - d_{ij(i)}, \mathbf{P}' \right) - \tilde{w}_k \left(j(i), I - d_{ij(i)}, \mathbf{P}' \right) \right] \tag{A7}
\end{aligned}$$

For $I \leq y_k^*(i, P) - 1$, the right-hand-side above reduces to the constant $-\tilde{P}_i$. For $I \geq y_k^*(i, P)$, the right-hand-side above reduces to $\tilde{w}_k \left(j(i), I + 1 - d_{ij(i)}, \mathbf{P}' \right) - \tilde{w}_k \left(j(i), I - d_{ij(i)}, \mathbf{P}' \right)$ which is non-decreasing by the convexity assumption on $\tilde{w}_k$. In order to establish convexity of $g(i, I, \mathbf{P})$, we are only left to check that

$$\tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w}_k \left(j(i), y_k^*(i, P) + 1 - d_{ij(i)}, \mathbf{P}' \right) - \tilde{w}_k \left(j(i), y_k^*(i, P) - d_{ij(i)}, \mathbf{P}' \right) \right] \geq -\tilde{P}_i,$$

which can be rewritten as

$$f_k(i, y_k^*(i, P) + 1, \mathbf{P}) - f_k(i, y_k^*(i, P), \mathbf{P}) \geq 0.$$

The latter is clearly true since f is minimized at $y_k^*(i, P)$. We conclude that $g(i, I, \mathbf{P})$ is a piecewise linear convex function with changes of slope contained in $B(i)$, which in turn implies that $\tilde{w}_{k+1}(i, I, \mathbf{P})$ is a piecewise linear convex function with changes of slope contained in $B(i)$.

We have hence established that the value iteration step preserves the convexity and piecewise linearity (with slope changes in $B(i)$) properties. This implies that the limit function $\tilde{w}$ is a piecewise linear convex function with respect to I , with changes of slope contained in $B(i)$.

Returning to Bellman's equation for the modified system, the optimal policy solves

$$\min_{q \in Q(i, j, I)} \left(\tilde{P}_i q + \tilde{\mathbb{E}}_{\mathbf{P}'|\mathbf{P}} \left[\tilde{w} \left(j(i), I + q - d_{ij(i)}, \mathbf{P}' \right) \right] \right) \tag{A8}$$

An argument similar to the one provided above establishes that the optimal policy is of the order-up-to type. There exist location and price dependent levels $S(i, \mathbf{P})$ such that $q^*(i, I, \mathbf{P}) = (S(i, \mathbf{P}) - I)^+$. In addition, those can be chosen so that $S(i, \mathbf{P}) \in B(i)$. This concludes the proof. $\square$

Proof of Proposition 2

Consider again the value iteration procedure detailed in (A1) in the proof of Proposition 1. Define

$$f_k(i, y, P_0, \epsilon, \alpha) := \tilde{P}_i y + \tilde{\mathbb{E}}_{P'_0|P_0, \epsilon', \alpha'|\alpha} [\tilde{w}_k(j(i), y - d_{ij(i)}, P'_0, \epsilon', \alpha')], \quad (\text{A9})$$

where P'_0 , ϵ' , and α' denote the crude oil price, the local price disturbance vector and the demand-supply state vector at the time the liner arrives at port $j^*(i)$. For all $0 \leq y \leq C - 1$, let

$$\Delta_k(i, y, P_0, \epsilon, \alpha) := f_k(i, y + 1, P_0, \epsilon, \alpha) - f_k(i, y, P_0, \epsilon, \alpha). \quad (\text{A10})$$

We have that

$$\tilde{w}_{k+1}(i, I, P_0, \epsilon, \alpha) = \min_{\max(I, d_{ij^*(i)}) \leq y \leq C} (f_k(i, y, P_0, \epsilon, \alpha)) - P_i I \quad (\text{A11})$$

Now, relax the lower bound constraint to $d_{ij^*(i)}$, and let $S_k(i, P_0, \epsilon, \alpha)$ be an optimal fuel-up-to level (i.e. the solution of $\min_{d_{ij^*(i)} \leq y \leq C} f_k(i, y, P_0, \epsilon, \alpha)$). Note that w_k is convex in I and hence so is $f_k(i, y, P_0, \epsilon, \alpha)$ as a function of y . This implies that at $S_k(i, P_0, \epsilon, \alpha)$, the following conditions need to be satisfied

$$\begin{aligned} \Delta_k(i, S_k(i, P_0, \epsilon, \alpha), P_0, \epsilon, \alpha) &\geq 0 \text{ if } d_{ij^*(i)} \leq S_k(i, P_0, \epsilon, \alpha) \leq C - 1 \\ \Delta_k(i, S_k(i, P_0, \epsilon, \alpha) - 1, P_0, \epsilon, \alpha) &\leq 0 \text{ if } d_{ij^*(i)} + 1 \leq S_k(i, P_0, \epsilon, \alpha) \leq C. \end{aligned} \quad (\text{A12})$$

Now, the fuel-up-to level $S_k(i, P_0, \epsilon, \alpha)$ is non-increasing in ϵ_i if $\Delta_k(i, y, P_0, \epsilon, \alpha)$ is non-decreasing in ϵ_i . The latter is clearly satisfied, since

$$\begin{aligned} \Delta_k(i, y, P_0, \epsilon, \alpha) &= \tilde{P}_i + \tilde{\mathbb{E}}_{P'_0|P_0, \epsilon', \alpha'|\alpha} [\tilde{w}_k(j(i), y - d_{ij(i)}, P'_0, \epsilon', \alpha') - \tilde{w}_k(j(i), y + 1 - d_{ij(i)}, P'_0, \epsilon', \alpha')] \\ &= (\gamma P_0 + \alpha_i + \epsilon_i) / \tau_{ij(i)} \\ &\quad + \tilde{\mathbb{E}}_{P'_0|P_0, \epsilon', \alpha'|\alpha} [\tilde{w}_k(j(i), y - d_{ij(i)}, P'_0, \epsilon', \alpha') - \tilde{w}_k(j(i), y + 1 - d_{ij(i)}, P'_0, \epsilon', \alpha')] \end{aligned} \quad (\text{A13})$$

and the second term does not depend on ϵ_i .

We now turn to the dependence on P_0 . The fuel-up-to level $S_k(i, P_0, \epsilon, \alpha)$ is non-increasing in P_0 if $\Delta_k(i, y, P_0, \epsilon, \alpha)$ is non-decreasing in P_0 . Suppose the relative value iteration algorithm is started with $\tilde{w}_0(\cdot) = 0$ for all states, then the condition is satisfied at stage $k = 0$. Suppose that it is also satisfied up to the stage k . At stage $k + 1$, we have for all $I \leq C - 1$,

$$\begin{aligned} w_{k+1}(i, I, P_0, \epsilon, \alpha) &= \tilde{P}_i (S_k(i, P_0, \epsilon, \alpha) - I)^+ + \tilde{\mathbb{E}}_{P'_0|P_0, \epsilon', \alpha'|\alpha} [w_k(j^*(i), I + (S_k(i, P_0, \epsilon, \alpha) - I)^+ - d_{ij^*(i)}, P'_0, \epsilon', \alpha')] \end{aligned} \quad (\text{A14})$$

and

$$\begin{aligned}
& \Delta_{k+1}(i, I, P_0, \epsilon, \alpha) \\
&= -\tilde{P}_i \mathbf{1}\{I \leq S_k(i, P_0, \epsilon, \alpha)\} \\
&\quad + \mathbf{1}\{I > S_k(i, P_0, \epsilon, \alpha)\} \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[\tilde{w}_k(j^*(i), I + (S_k(i, P_0, \epsilon, \alpha) - I)^+ - d_{ij^*(i)} + 1, P'_0, \epsilon', \alpha') \right. \\
&\quad \left. - \tilde{w}_k(j^*(i), I + (S_k(i, P_0, \epsilon, \alpha) - I)^+ - d_{ij^*(i)}, P'_0, \epsilon', \alpha') \right], \\
&= -\tilde{P}_i \mathbf{1}\{I \leq S_k(i, P_0, \epsilon, \alpha)\} + \mathbf{1}\{I > S_k(i, P_0, \epsilon, \alpha)\} \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[\Delta_k(i, I - d_{ij^*(i)}, P_0, \epsilon, \alpha) \right] \\
&= -\tilde{P}_i + \mathbf{1}\{I > S_k(i, P_0, \epsilon, \alpha)\} \left[\tilde{P}_i + \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[\Delta_k(i, I - d_{ij^*(i)}, P_0, \epsilon, \alpha) \right] \right]. \tag{A15}
\end{aligned}$$

Let $j = j^*(i)$ and $k = j^*(j^*(i))$, $P_0 = P_0(t)$, $P'_0 = P_0(t + \tau_{ij})$ and $P''_0 = P_0(t + \tau_{ij} + \tau_{jk})$. Using the (A15), we have

$$\begin{aligned}
& \tilde{P}_i + \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[\Delta_{k+1}(i, I - d_{ij^*(i)}, P_0, \epsilon, \alpha) \right] \\
&= \tilde{P}_i + \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[-\tilde{P}'_j + \mathbf{1}\{I > S_k(i, P_0, \epsilon, \alpha)\} \left[-\tilde{P}'_k + \mathbb{E}_{P''_0|P'_0, \epsilon'', \alpha''|\alpha'} \left[\Delta_k(i, I - d_{jk}, P'_0, \epsilon, \alpha') \right] \right] \right] \\
&= \gamma P_0 + \alpha_i + \epsilon_i \\
&\quad + \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[-\gamma P'_0 - \alpha'_j - \epsilon'_j + \mathbf{1}\{I > S_k(j, P'_0, \epsilon', \alpha')\} \left[\tilde{P}'_k + \mathbb{E}_{P''_0|P'_0, \epsilon'', \alpha''|\alpha'} \left[\Delta_k(i, I - d_{jk}, P'_0, \epsilon, \alpha') \right] \right] \right] \\
&= A + B,
\end{aligned}$$

where

$$\begin{aligned}
A &= \gamma[P_0 - \mathbb{E}_{P'_0|P_0}[P'_0]] + \alpha_i - \mathbb{E}_{\alpha'|\alpha}[\alpha'_j] + \epsilon_i, \\
B &= \mathbb{E}_{P'_0|P_0, \epsilon', \alpha'|\alpha} \left[\mathbf{1}\{I > S_k(j, P'_0, \epsilon', \alpha')\} \left[\tilde{P}'_k + \mathbb{E}_{P''_0|P'_0, \epsilon'', \alpha''|\alpha'} \left[\Delta_k(i, I - d_{jk}, P'_0, \epsilon, \alpha') \right] \right] \right]. \tag{A16}
\end{aligned}$$

We first analyze the term B . $\tilde{P}'_k + \mathbb{E}_{P''_0|P'_0, \epsilon'', \alpha''|\alpha'} \left[\Delta_k(i, I - d_{jk}, P'_0, \epsilon, \alpha') \right]$ is non-decreasing in P'_0 by the induction hypothesis and note that this expression is non-negative on the domain $\{I > S_k(i, P_0, \mathbf{P}, \alpha)\}$ (by the optimality condition). The induction hypothesis also implies that $S_k(i, P'_0, \epsilon', \alpha')$ is non-increasing in P'_0 and hence $\mathbf{1}\{I > S_k(i, P'_0, \epsilon', \alpha')\}$ is non-decreasing in P_0 . We deduce that the product term within the outer expectation is non-decreasing in P'_0 . Assumption 1 (i) (the stochastic ordering of prices) now implies that B is non-decreasing in P_0 .

Turning to A , one notes that the term $\gamma[P_0 - \mathbb{E}_{P'_0|P_0}[P'_0]]$ is non-decreasing in P_0 by Assumption 1 (ii). We deduce that $\Delta_{k+1}(i, y, P_0, \epsilon, \alpha)$ is non-decreasing in P_0 . This in turn implies that the property is satisfied in the limit and the result follows. $\square$

Proof of Proposition 3

Let $I' > I$, both in $[0, C]$. By definition of $q^*(i, j, I, \mathbf{P})$, we have

$$q^*(i, j, I, \mathbf{P}) = -I + \min \left(\arg \max_{y \in \tilde{Q}(i, j, I)} \left(-P_i q + \mathbb{E}_{\mathbf{P}'|\mathbf{P}, \mathbf{r}'} \left[h \left(j, y - d_{ij}, \mathbf{P}', \mathbf{r}' \right) \right] \right) \right),$$

where $\tilde{Q}(i, j, I) = \{y | y \geq I, d_{ij} \leq y \leq C\}$ and y is the fuel-up-to level. Now, we have

$$\begin{aligned}
& I + q^*(i, j, I, \mathbf{P}) \\
&= \min \left(\arg \max_{y \in \tilde{Q}(i, j, I)} \left(-P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}, \mathbf{r}'} \left[h(j, y - d_{ij}, \mathbf{P}', \mathbf{r}') \right] \right) \right) \\
&\leq \min \left(\arg \max_{y \in \tilde{Q}(i, j, I')} \left(-P_i q + \mathbb{E}_{\mathbf{P}' | \mathbf{P}, \mathbf{r}'} \left[h(j, y - d_{ij}, \mathbf{P}', \mathbf{r}') \right] \right) \right) \\
&= I' + q^*(i, j, I', \mathbf{P}),
\end{aligned} \tag{A17}$$

where the inequality follows from the fact that $I' > I$. Hence, $I + q_k^*(i, j, I, \mathbf{P})$ is a non-decreasing function of I . $\square$

Proof of Proposition 4

a) Consider without loss of generality the point $z_1 = (i_1, I_1)$. Recall that, according to our assumptions, i_1 is reachable from i in a finite number transitions. Consider a vessel making such transitions while buying the minimum possible amount of fuel necessary for such a trip. Let $I_1^{(0)}$ denote the level of fuel inventory that a vessel has when arriving at i_1 . Note that such inventory level need not be equal to I_1 . If $I_1^{(0)} \leq I_1$, then moving from $(i_1, I_1^{(0)})$ to $z_2 = (i_2, I_2)$ in a single transition is feasible, since $I_1 - d_{i_1 i_2} \leq I_2 \leq C - d_{i_1 i_2}$ implies $I_1^{(0)} - d_{i_1 i_2} \leq I_2 \leq C - d_{i_1 i_2}$. Now, one can move from (i_2, I_2) to (i_1, I_1) in a finite number of steps by following the cycle and the result follows. Suppose, on the other hand, that $I_1^{(0)} > I_1$. Observe that on the cycle, the purchased quantity of fuel at the node z_l is given by $q_{i_l, i_{l+1}} = I_{l+1} + d_{i_l, i_{l+1}} - I_l$, $l = 1, \dots, k$, where $i_{k+1} = i_1$, $I_{k+1} = I_1$. Let m^* be the maximum integer $m \geq 1$ such that $I_1^{(0)} - \sum_{l=1}^{m-1} d_{i_l, i_{l+1}} > I_m$, where i_n for $n > k$ is defined as $i_{n \bmod k}$. Starting at $(i_1, I_1^{(0)})$, consider the path that purchases no fuel until arriving at i_{m^*} . Such policy is feasible by definition of m^* . Then, at i_{m^*} , the inventory $I_m^{(0)}$ satisfies $I_m^{(0)} \leq I_{m^*+1} - d_{i_{m^*}, i_{m^*+1}}$. While i_{m^*} , consider the purchase of $I_{m^*+1} - d_{i_{m^*}, i_{m^*+1}} - I_m^{(0)} \geq 0$ units of fuel. Then a vessel arrives at (i_{m^*+1}, I_{m^*+1}) and can now reach (i_1, I_1) in a finite number of steps by following the cycle.

b) This part just indicates how to reach the optimal cycle in the location-inventory space.

$\square$

Proof of Proposition 5

Consider the revenue maximizing cycle $\mathcal{C}^{(r)}$. Let

$$k = \left\lfloor (C - d_{i_{m^*}, i_{\min}} - d_{i_{m^*}, i_{\min}}) / \sum_{i=1}^m d_{i, i+1} \right\rfloor,$$

and note that $k \geq 1$ since $C \geq I_{\mathcal{C}^{(r)}}$. The associated profit per unit of time corresponding to $\pi_{\mathcal{C}^{(r)}}$ is given by

$$\lambda_{\pi_{\mathcal{C}^{(r)}}} = \frac{k \sum_{(i,j) \in \mathcal{C}^{(r)}} r_{ij} - P_{\min} \left[k \sum_{(i,j) \in \mathcal{C}^{(r)}} d_{ij} + d_{i_{m^*}, i_{\min}} + d_{i_{\min}, i_{m^*}} \right]}{k \sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij} + \tau_{i_{m^*}, i_{\min}} + \tau_{i_{\min}, i_{m^*}}}$$

$$= \frac{\sum_{(i,j) \in \mathcal{C}^{(r)}} [r_{ij} - P_{\min} d_{ij}] + k^{-1} P_{\min} [d_{i_m^*, i_{\min}} + d_{i_{\min}, i_m^*}]}{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij} + k^{-1} [\tau_{i_m^*, i_{\min}} + \tau_{i_{\min}, i_m^*}]}$$

Note that λ^D is upper bounded by the optimal profits in a system where all prices are set at $P_{\min}$, and that in such a system, the vessel maximizes profits by following $\mathcal{C}^{(r)}$ (see Section 6.2.1). Hence

$$\begin{aligned} \frac{\lambda_{\pi_{\mathcal{C}^{(r)}}}}{\lambda^D} &\geq \frac{\sum_{(i,j) \in \mathcal{C}^{(r)}} [r_{ij} - P_{\min} d_{ij}] + k^{-1} P_{\min} [d_{i_m^*, i_{\min}} + d_{i_{\min}, i_m^*}]}{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij} + k^{-1} [\tau_{i_m^*, i_{\min}} + \tau_{i_{\min}, i_m^*}]} \frac{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij}}{\sum_{(i,j) \in \mathcal{C}^{(r)}} [r_{ij} - P_{\min} d_{ij}]} \\ &= \frac{\sum_{(i,j) \in \mathcal{C}^{(r)}} [r_{ij} - P_{\min} d_{ij}] + k^{-1} P_{\min} [d_{i_m^*, i_{\min}} + d_{i_{\min}, i_m^*}]}{\sum_{(i,j) \in \mathcal{C}^{(r)}} [r_{ij} - P_{\min} d_{ij}]} \frac{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij}}{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij} + k^{-1} [\tau_{i_m^*, i_{\min}} + \tau_{i_{\min}, i_m^*}]} \\ &\geq \left[1 + k^{-1} \frac{\tau_{i^*l} + \tau_{i^*l}}{\sum_{(i,j) \in \mathcal{C}^{(r)}} \tau_{ij}} \right]^{-1} \\ &\geq 1 - k^{-1} \frac{d_{i_m^*, i_{\min}} + d_{i_{\min}, i_m^*}}{\sum_{(i,j) \in \mathcal{C}^{(r)}} d_{ij}} \end{aligned}$$

Note that $C \geq k \sum_{(i,j) \in \mathcal{C}^{(r)}} d_{ij} + [d_{i_m^*, i_{\min}} + d_{i_{\min}, i_m^*}]$ and hence $(1/C)k \sum_{(i,j) \in \mathcal{C}^{(r)}} d_{ij} \leq 1 - 2\zeta$.

We deduce that

$$\frac{\lambda_{\pi_{\mathcal{C}^{(r)}}}}{\lambda^D} \geq 1 - \frac{2\zeta}{1 - 2\zeta} = \frac{1 - 4\zeta}{1 - 2\zeta}.$$

This concludes the proof. $\square$

Proof of Proposition 6

Let $\hat{\lambda}$ denote the optimal profit per unit of time in a system where for all $i \in \mathcal{N}$, the price at port i is constant equal to $\underline{P}$ and for all arcs $(i, j) \in \mathcal{N}$, the rewards are constant and equal to $\bar{r}_{ij}$. Let $\hat{\mathcal{C}}$ denote the optimal (location) cycle corresponding to $\hat{\lambda}$. We have $\hat{\lambda} \geq \lambda^*$ implying that $\lambda^* - \lambda_{OLC} \leq \hat{\lambda} - \lambda_{OLC}$. Below, we upper bound the $\hat{\lambda} - \lambda_{OLC}$.

Consider a system where all the fuel prices are set at $P_c = \underline{P}$ and rewards are set at their steady-state expectations and note in such a system, the optimal cycle to follow is the one followed by the OLC heuristic. Let us denote this cycle by $\mathcal{C}_{OLC}$. In particular, the profit per unit of time associated with this cycle in this system is no less than the profit per unit of time generated when following $\hat{\mathcal{C}}$ and

$$\begin{aligned} \frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} \mathbb{E}[r_{ij}] - P_c d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}} &\geq \frac{\sum_{(i,j) \in \hat{\mathcal{C}}} \mathbb{E}[r_{ij}] - P_c d_{ij}}{\sum_{(i,j) \in \hat{\mathcal{C}}} \tau_{ij}} \\ &= \frac{\sum_{(i,j) \in \hat{\mathcal{C}}} \bar{r}_{ij} - \underline{P} d_{ij}}{\sum_{(i,j) \in \hat{\mathcal{C}}} \tau_{ij}} - \frac{\sum_{(i,j) \in \hat{\mathcal{C}}} (\bar{r}_{ij} - \mathbb{E}[r_{ij}])}{\sum_{(i,j) \in \hat{\mathcal{C}}} \tau_{ij}} \\ &= \hat{\lambda} - \frac{\sum_{(i,j) \in \hat{\mathcal{C}}} (\bar{r}_{ij} - \mathbb{E}[r_{ij}])}{\sum_{(i,j) \in \hat{\mathcal{C}}} \tau_{ij}} \\ &\geq \hat{\lambda} - \frac{\sum_{(i,j) \in \hat{\mathcal{C}}} 1}{\sum_{(i,j) \in \hat{\mathcal{C}}} \tau_{ij}} \Delta_r. \end{aligned}$$

Now, in the true system, consider the policy that follows the cycle $\mathcal{C}_{OLC}$ and replenishes at each port the exact amount necessary to reach the next port on the cycle. Its performance is given by

$$\frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} \mathbb{E}[r_{ij}] - \mathbb{E}_\infty[P_i]d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}},$$

and is dominated by the performance of the OLC heuristic as the latter follows the same route but optimizes the refueling amounts. Hence

$$\begin{aligned} \lambda_{OLC} &\geq \frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} \mathbb{E}[r_{ij}] - \mathbb{E}_\infty[P_i]d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}} \\ &= \frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} \mathbb{E}[r_{ij}] - \underline{P}d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}} - \frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} (\mathbb{E}_\infty[P_i] - \underline{P})d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}} \\ &\geq \hat{\lambda} - \omega\Delta_r - \frac{\sum_{(i,j) \in \mathcal{C}_{OLC}} (\mathbb{E}_\infty[P_i] - \underline{P})d_{ij}}{\sum_{(i,j) \in \mathcal{C}_{OLC}} \tau_{ij}} \\ &\geq \hat{\lambda} - \omega\Delta_r - \rho\Delta_p. \end{aligned}$$

This completes the proof. $\square$

Proof of Proposition 7

Let $\hat{\lambda}$ denote the optimal profit per unit of time in a system where for all $i \in \mathcal{N}$, the price at port i is constant equal to $\underline{P}_i$ and for all arcs $(i, j) \in \mathcal{N}$, the rewards are constant and equal to $\bar{r}_{ij}$. Let $\hat{\mathcal{C}}$ denote the optimal (location-inventory) cycle corresponding to $\hat{\lambda}$. We have $\hat{\lambda} \geq \lambda^*$ implying that $\lambda^* - \lambda_{OLIC} \leq \hat{\lambda} - \lambda_{OLIC}$. Below, we upper bound the $\hat{\lambda} - \lambda_{OLIC}$.

Consider a system where all the fuel prices and rewards are set at their steady-state expectations and note in such a system, λ_{OLIC} is the optimal profit per unit of time. In particular, λ_{OLIC} is no less than the long-run profit per unit of time generated when following $\hat{\mathcal{C}}$ and

$$\begin{aligned} \lambda_{OLIC} &\geq \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} \mathbb{E}[r_{zz'}] - \mathbb{E}_\infty[P_z](I' - I + d_{zz'})}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} \\ &= \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} \bar{r}_{zz'} - \underline{P}_z(I' - I + d_{zz'})}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} - \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} (\bar{r}_{zz'} - \mathbb{E}[r_{zz'}]) + (\mathbb{E}_\infty[P_z] - \underline{P}_z)(I' - I + d_{zz'})}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} \\ &= \hat{\lambda} - \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} (\bar{r}_{zz'} - \mathbb{E}[r_{zz'}]) + (\mathbb{E}_\infty[P_z] - \underline{P}_z)(I' - I + d_{zz'})}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} \\ &\geq \hat{\lambda} - \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} 1}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} \Delta_r - \frac{\sum_{(z,z') \in \hat{\mathcal{C}}} d_{zz'}}{\sum_{(z,z') \in \hat{\mathcal{C}}} \tau_{zz'}} \tilde{\Delta}_p \\ &\geq \hat{\lambda} - \tilde{\omega}\Delta_r - \tilde{\rho}\tilde{\Delta}_p. \end{aligned}$$

This yields the result. $\square$

B Estimation Details of the Fuel Pricing Model

Below we describe how we used the actual pricing data to estimate the parameters associated with the price model (3). Let k denote the number of price observations (days) and let P_{ij}^k denote the price at port $i = 1, \dots, N$ on day $j = 1, \dots, k$. The corresponding price of crude oil on day $j = 1, \dots, k$ is denoted as P_{0j}^k . Define

$$\hat{\gamma}^k = \frac{\sum_{i=1}^N \sum_{j=1}^k P_{ij}^k}{N \sum_{j=1}^k P_{0j}^k}. \quad (\text{B1})$$

Note that in our model we assume that $\sum_{i=1}^N A_i = 0$, so that we get $\hat{\gamma}^k \rightarrow \gamma$ as $k \rightarrow \infty$ almost surely. Similarly, define

$$\hat{A}_i^k = \frac{\sum_{j=1}^k P_{ij}^k}{k} - \hat{\gamma}^k \frac{\sum_{j=1}^k P_{0j}^k}{k}, i = 1, \dots, N, \quad (\text{B2})$$

and note that $\hat{A}_i^k \rightarrow A_i$ as $k \rightarrow \infty$ almost surely. Further, let

$$\bar{P}_{ij}^k = P_{ij}^k - \hat{\gamma}^k P_{0j}^k - \hat{A}_i^k, i = 1, \dots, N, j = 1, \dots, k \quad (\text{B3})$$

denote the deviation of the fuel price at port $i = 1, \dots, N$ on day $j = 1, \dots, k$ from the “average” behavior. If the fuel price dynamics is described by the model (3), the “trajectory” of $\bar{P}_{ij}^k$ is a result of juxtaposition of two independent, mean-zero processes: a two-state Markov process with amplitude δ_i and inertia parameter η_i and the stationary process with standard deviation σ_i . Given the state information for the port-specific Markov chains were not available to us, we estimated the remaining price dynamics parameters employing the following simple approach. Assuming that σ_i is substantially smaller than δ_i , we can estimate the value of η_i as the fraction of days on which the sign of $\bar{P}_{ij}^k$ coincides with the sign of $\bar{P}_{i,j-1}^k$:

$$\hat{\eta}_i^k = \frac{\sum_{j=2}^k \left(\frac{\bar{P}_{ij}^k \bar{P}_{i,j-1}^k + |\bar{P}_{ij}^k \bar{P}_{i,j-1}^k|}{2 |\bar{P}_{ij}^k \bar{P}_{i,j-1}^k|} \right)}{k-1}, i = 1, \dots, N. \quad (\text{B4})$$

The adequacy of this assumption is confirmed post-factum by the values shown in Table 1. In turn, the values of δ_i and σ_i are estimated as

$$\hat{\delta}_i^k = \frac{\sum_{j=1}^k |\bar{P}_{ij}^k|}{k}, i = 1, \dots, N, \quad (\text{B5})$$

and

$$\hat{\sigma}_i^k = \sqrt{\frac{\sum_{j=1}^k \left(|\bar{P}_{ij}^k| - \hat{\delta}_i^k \right)^2}{k-1}}, i = 1, \dots, N. \quad (\text{B6})$$

In our numerical analysis we use the parameter estimates obtained from the prices for $N = 18$ ports for the period of January-June 2005 ($k = 98$ trading days). Thus, in total, $18 \times 98 = 1764$ ($N \times k$) pricing data points were used to estimate $4 \times N + 1 = 4 \times 18 + 1 = 73$ parameters of our bunkers pricing model.

Due to the absence of the port-specific information on the realized states of the Markov chains α , we could not provide reliable confidence intervals for the parameter estimates in Table 1. However, in order to get a sense of the goodness-of-fit for our model, we have also calculated the resulting R^2 value for our model as well as R^2 value for a simpler model of bunkers price dynamics, which ignored the time dynamics of the local supply corrections to bunkers prices, replacing them by port-dependent constants ($P_i(t) = \gamma P_0(t) + A_i + \epsilon_i(t)$, where A_i is constant). The R^2 for our model turned out to be 0.68 as compared to 0.35 for the simplified model. Clearly, the presence of the Markov term α accounts for a significant fraction of the explanatory power of our model. As the value $\gamma = 4.43$ indicates, the oil-price-related component $\gamma P_0(t)$ is responsible for the bulk of the bunkers price value: for the average crude oil price over the period of January-June 2005, \$51.4, its contribution is \$227.7. On a given day at a given port, this contribution accounts for about 80% – 90% of the total bunker price.

We now turn attention to the estimation of the dynamics of the crude oil prices. Below we use the actual WTI crude oil prices for the period of January-June 2005 to fit both models described in (4) and (5). Figure B1 shows the histogram of actual relative daily changes in oil prices for this period along with the plot of the PDF of the corresponding best-fit normal random variable ($\mu_0 = 0.00$, $\sigma_0 = 0.03$). The goodness-of-fit test measures for this model are: Anderson-Darling, $AD = 0.8089$, Kolmogorov-Smirnov, $KS = 0.1106$, and Chi-Square, $\chi^2 = 27.04$.

The fit for the mean-reverting model,

$$P_0(t+1) = \lambda P_0(t) + (1-\lambda)\hat{P}_0 + \varepsilon_0(t), \quad \varepsilon_0(t) \sim N(0, \sigma_0), \quad (\text{B7})$$

with $\lambda = 0.96$, $\hat{P}_0 = \$55.5$, and $\sigma_0 = \$1.25$, illustrated in Figure B2, is near-perfect ($R^2 = 0.91$, significance of $F = 0$, 95% CI for the slope value λ is $[0.90, 1.02]$). The goodness-of-fit values measures for the normal distribution of error terms for this model ($AD = 0.7076$, $KS = 0.1055$, $\chi^2 = 20.75$) indicate a somewhat better fit than that for the GBM model.

Based on these estimation results, we have selected the mean-reverting model to represent stochastic crude oil price dynamics. In addition to the goodness-of-fit arguments, computational arguments also favor selecting mean-reverting dynamics over those of a GBM model since the former allow

Figure B1: Histograms of actual and best-fit normal relative daily crude oil price changes for January-June 2005.

Figure B2: Actual and best-fit mean-reverting values of the crude-oil prices for January-June 2005.

the use of a finite number of prices for describing the state of the system in the dynamic program (2). Indeed, since the limiting variance of the mean-reverting process (B7),

$$\lim_{t \rightarrow \infty} (\text{Var}[P_0(t)]) = \frac{(1.25)^2}{1 - (0.96)^2} = 19.92, \quad (\text{B8})$$

is finite (as opposed to the corresponding value for the GBM process), we can use the finite-state approximation to represent this process. Consider the “limiting” value of the mean-reverting process (B7), $\lim_{t \rightarrow \infty} (\mathbb{E}[P_0(t)]) = 55.5$. Since the limiting value of the standard deviation of $P_0(t)$ from this level, $\sqrt{19.92} = 4.46$, we can limit our attention to the “2 deviations” price interval $[55.5 - 2 \times 4.46, 55.5 + 2 \times 4.46]$, or, rounding, to $[46, 62]$ ⁵. Noting that since according to (B7), the conditional distribution of $P_0(t+1)$ is normal with mean $0.96P_0(t) + (1 - 0.96) \times 55.5$ and standard deviation of $\sigma_0 = 1.25$, we can use a standard discretization procedure to approximate this continuous distribution by a discrete one defined on the interval $[P_0^{\min}, P_0^{\max}]$ with $P_0^{\min} = 46$ and $P_0^{\max} = 62$ as follows. Letting, for any P , $f_+(P) = [P + 0.5 - (\lambda P_0(t) + (1 - \lambda) \hat{P}_0)]\sigma_0^{-1}$ and $f_-(P) = [P - 0.5 - (\lambda P_0(t) + (1 - \lambda) \hat{P}_0)]\sigma_0^{-1}$, we define

$$\text{Prob}(P_0(t+1) = P | P_0(t)) = \begin{cases} 1 - \Phi(f_-(P)), & P = P_0^{\max}, \\ \Phi(f_+(P)) - \Phi(f_-(P)), & P_0^{\min} < P < P_0^{\max}, \\ \Phi(f_+(P)), & P = P_0^{\min}, \end{cases} \quad (\text{B9})$$

where $\lambda = 0.96$, $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt$ is the standard normal CDF, and $P_0(t) \in [P_0^{\min}, P_0^{\max}]$.

C Numerical Study of the Liner Case: Fuel-up-to Levels and Heuristic Policies.

Below we report the results of a numerical study designed to provide insights on the influence of various problem parameters as well as to test the performance of possible heuristics. We start by introducing the data that pertains to the network of ports and to the liner vessel.

C.1 Network and Vessel Data

We consider an example of a liner with fuel capacity of 3000 metric tons (mts) that operates between the following ports: Santa Marta in Colombia, Moin in Costa Rica, Castilla in Honduras, Port Everglades in Florida and Wilmington in Delaware. According to the time which has to be spent in each port (which includes, e.g., the time associated with the loading/unloading activities),

⁵In May of 2008, the crude oil has crossed a \$120/bbl barrier. This development underscores a precarious nature of the process of forecasting the crude oil prices, and suggests frequent recalibration of the forecasting models.

the ship can only replenish fuel at Wilmington (port 1) or Moin (port 2). The approximate fuel consumption for the trip from Wilmington to Moin and from Moin to Wilmington are given by 400 mts and 600 mts, respectively, and the time associated with any of the two trips is approximately 5 days. Note that the two one-way consumption values differ as the respective one-way routes include different stops.

Based on the above description, we model the network for the liner under consideration as consisting of $N = 2$ ports, with the following parameters: $\tau_{12} = \tau_{21} = 5$, $d_{12} = 400$ mts and $d_{21} = 600$ mts. To simplify computations, we assume that the liner can only replenish fuel by multiples of 200 mts. In our numerical study we use (3) as the model bunkers price dynamics with parameters estimated using (B1)-(B6). The base-case values of the price parameters for the ports of Wilmington and Moin are shown in Table B1.

Parameter	Wilmington	Moin
A , \$/mts	0.0	-1.7
δ , \$/mts	29.2	25.8
η	0.99	0.97
σ , \$/mts	14.4	14.3

Table B1: Parameters of the bunkers fuel price dynamics for the ports of Wilmington and Moin.

We use a simple form for random price correction terms by assuming that ε_i at port i takes two values, $\pm\sigma_i$, w.p. 0.5 for each value. The oil price dynamics is taken to follow the mean-reverting model (B9) with $\sigma_0 = 1.25$, $\lambda = 0.96$, $P_0^{\min} = 46$ and $P_0^{\max} = 62$. In summary, the state space of the resulting dynamic program contains 4352 states.

For the base-case values of problem parameters, the optimal value of the expected daily fuel cost, obtained by solving the DP (6), is \$21,400 per day. As an illustration of the optimal refueling policy, Figure B1 shows the optimal fuel-up-to-levels as functions of the oil price in the “low price corrections” states with $\alpha_i = A_i - \delta_i$, $\varepsilon_i = -\sigma_i$, $i = \text{WIL, MOI}$.

In this case, MOI is the more expensive port out of the two for any oil price level, and the refuel-up-to levels for MOI do not exceed those for WIL. In addition, those levels naturally exhibit drops from “fuel-up-to-capacity” at both ports when the oil price is at its lowest ($P_0 = 46$) to “single roundtrip” refueling at either port when the oil price cannot be any higher ($P_0 = 62$). Note that even for the highest oil price level the negative price corrections $\varepsilon_i = -\sigma_i$, $i = \text{WIL, MOI}$ make it beneficial to move from “refuel as little as possible” to “single roundtrip” decisions at both ports.

Figure B1: Base-case scenario: critical refueling indices $S(i, P_0)$ for $i = \text{WIL}, \text{MOI}$, and $P_0 = 46, \dots, 62$ ($\alpha_i = A_i - \delta_i$, $\varepsilon_i = -\sigma_i$, $i = \text{WIL}, \text{MOI}$).

C.2 Fuel-up-to Levels: Sensitivity Properties

In order to understand how the refueling decisions depend on the parameters of the bunkers price dynamics, we would like to deviate from the base-case by exploring a wider range of problem settings. In particular, we first consider a variant of our problem for which the bunkers prices at both ports at any point in time are completely determined by the value of the oil price at that point in time (i.e., the case with $\delta_{\text{WIL}} = \delta_{\text{MOI}} = \sigma_{\text{WIL}} = \sigma_{\text{MOI}} = 0$, while the rest of the parameters are set at their base-case values). The resulting values of the critical refueling indices $S(i, P_0)$ for $i = \text{WIL}$ and $i = \text{MOI}$ and for $P_0 = [46, 62]$ are shown in Figure B2. In this case, opposite to what we observed in the base-case setting, WIL is the more expensive port out of the two for any price level, and the refuel-up-to levels for WIL do not exceed those for MOI. As in the base scenario, those levels exhibit sharp drops from “fuel-up-to-capacity” at both ports when the oil price is at its lowest to “refuel as little as possible” at either port when the oil price is at its highest.

Figure B3 shows how the critical refueling indices observed in the previous “zero-case” scenario change when the daily adjustment factor ε_{WIL} is introduced (with the base-case value $\sigma_{\text{WIL}} = 14.4$). Note that the degeneracy of the base-case $S(\text{WIL})$ curve from Figure B2 is lifted and the refueling indices for WIL “split” into two curves, one below (corresponding to $\varepsilon_{\text{WIL}} = 14.4$) and one above (corresponding to $\varepsilon_{\text{WIL}} = -14.4$) the base-case curve. Such split is expected in the view of higher-than-average (lower-than-average) prices in WIL corresponding to the lower (higher) curve. Another interesting feature depicted on Figure B3 is “lowering” of the critical index curve for MOI:

Figure B2: Critical refueling indices $S(i, P_0)$ for $i = \text{WIL}, \text{MOI}$ and $P_0 = 46, \dots, 62$ ($\delta_i = \sigma_i = 0$, $i = \text{WIL}, \text{MOI}$).

for example, when $P_0 = 51$, it is optimal when in MOI to fuel-up-to the level corresponding to one round trip instead of up to vessel's capacity. Such a change reflects the decreased attractiveness of MOI as a port of purchase in presence of newly-found occasional bargain at WIL whenever $\varepsilon_{\text{WIL}} = -14.4$.

Figure B4 probes the effects of the presence of the supply correction at WIL. In this example the amplitude δ_{WIL} of the WIL supply correction is set at the base-case value of 29.2. Note that an important difference between the critical index curves in this case and those of Figure B3 is the splitting off of the MOI curve corresponding to the states with “high” WIL prices from the curve corresponding the “low” WIL prices. While in the case of random price corrections depicted in Figure B3 the value of correction observed at WIL does not influence refueling decision made at the same time in MOI, this is not the case when such price corrections follow high-inertia Markov pattern: a “high” price observed at WIL at the time of MOI refueling decision has 99% chance of remaining “high” when the vessel arrives at WIL, thus raising the refueling level at MOI, as in the case of $P_0 = 53$. Similarly, a “low” WIL price has 99% chance of remaining “high” when the vessel arrives at WIL, thus lowering the refueling level at MOI, as in the case of $P_0 = 53$. Note that this argument remains valid only if the inertia coefficient η_{WIL} is above 0.5 - i.e., if observed “high” price makes the “high” price more likely to be observed next day. When η_{WIL} drops below 0.5, the argument is reversed: “high” price today makes “low” price more likely tomorrow. This case is illustrated in Figure B5 where $\eta_{\text{WIL}} = 0.01$: the MOI curves corresponding to “high” and

Figure B3: Critical refueling indices $S(i, P_0)$ for $i = \text{WIL}, \text{MOI}$, $\sigma_{\text{WIL}} = 14.4$, and $P_0 = 46, \dots, 62$ ($\delta_{\text{WIL}} = \delta_{\text{MOI}} = \sigma_{\text{MOI}} = 0$).

“low” WIL prices are, naturally, reversed as compared to Figure B4.

C.3 Numerical Results on Heuristic Policies

In our numerical study we also investigated the impact of the variations in the values of the local supply correction parameters, the local supply inertias and daily adjustment factor volatilities on the optimal cost performance. In addition we tested the performance of several heuristics that could serve as alternatives to the optimal policy in cases where the dynamic program (6) is too large to solve to optimality. For that purpose, we computed the long-run average daily cost μ^π associated with each of the following five policies⁶:

1. μ^{opt} - the value of the optimal policy.
2. μ^{M} - the value of the policy defined as follows: whenever in Wilmington, buy up to 400 mts and whenever in Moin, buy up to 600 mts. Note that this policy only buys the fuel necessary to get to the next location; in that sense, it is fully myopic.
3. μ^{MW} - the value of the policy defined as follows: whenever in Wilmington, buy up to 1000 mts and whenever in Moin, buy up to 600 mts. Note that when using this policy, the shipping

⁶Note that while in the previous section we have analyzed the properties of the finite-horizon problem, here we focus on the long-run expected cost per day, as this profit measure is independent of initial conditions.

Figure B4: Critical refueling indices $S(i, P_0, \alpha_{WIL})$ for $i = WIL, MOI$, $\delta_{WIL} = 29.2$, and $P_0 = 46, \dots, 62$ ($\sigma_{WIL} = \delta_{MOI} = \sigma_{MOI} = 0$).

Figure B5: Critical refueling indices $S(i, P_0, \alpha_{rmWIL})$ for $i = WIL, MOI$, $\delta_{WIL} = 29.2$, $\eta_{WIL} = 0.01$, and $P_0 = 46, \dots, 62$ ($\sigma_{WIL} = \delta_{MOI} = \sigma_{MOI} = 0$).

company effectively only purchases fuel at Wilmington after the first cycle.

4. μ^{MM} - the value of the policy defined as follows: whenever in Wilmington, buy up to 400 mts and whenever in Moin, buy up to 1000 mts. Note that when using this policy, the shipping company, after the first cycle, only purchases fuel at Moin.
5. μ^{CE} - the value of the “certainty equivalence” policy: at each stage, the policy is computed assuming that the prices will stay constant at their current levels for all future periods.

For each policy π , we define the associated optimality ratio as: $\mathcal{R}^\pi = (\mu^\pi - \mu^{\text{opt}})/\mu^{\text{opt}}$. Note that currently, the marine shipping company always purchases fuel in Wilmington and thus, essentially, follows the policy $\pi = \text{MW}$.

First, we have computed the optimality ratios for four heuristic policies for the base-case scenario described in Section 4 and Table B1. The “certainty equivalence” heuristic CE exhibits the best performance with the daily expected cost of \$22,397 per day (as compared with the optimal value of \$21,400/day) and the optimality ratio of $\mathcal{R}^{\text{CE}} = 4.66\%$. Other heuristics show the following performances: $\mathcal{R}^{\text{M}} = 13.68\%$, $\mathcal{R}^{\text{MW}} = 13.87\%$, $\mathcal{R}^{\text{MM}} = 13.56\%$. Strong performance of the CE heuristic is helped by a high values of η_{WIL} and η_{MOI} and, thus, by high inertia exhibited by differences in prices in Wilmington and Moin. In the remainder of the Section we investigate the effects of some of parameters of the bunkers price dynamics on the performance of the above heuristic policies.

The effect of the supply correction amplitude δ_{WIL} . In order to isolate the effects of the corrections due to supply availability, we first assume that $\sigma_{\text{WIL}} = \sigma_{\text{MOI}} = 0$, i.e., that there is no daily random correction in addition to the supply correction. Also, we assume that $\delta_{\text{MOI}} = 0$ and analyze the performance of the policies outlined above as δ_{WIL} is varied. Throughout, η_{WIL} is set to 0.99 reflecting a base-case level inertia in the balance of demand and supply factors at Wilmington. Note that in this setting, the fuel prices at Moin are exclusively driven by the oil price P_0 while the prices at Wilmington are driven by both P_0 and the supply correction α_{WIL} .

δ_{WIL}	$R^{\text{M}}(\%)$	$R^{\text{MW}}(\%)$	$R^{\text{MM}}(\%)$	$R^{\text{CE}}(\%)$
0	2.87	3.30	2.58	2.58
10	4.36	4.71	4.13	2.23
20	6.73	7.00	6.56	2.32
30	9.40	9.57	9.28	2.54

Table B2: Relative performance of heuristics as a function of δ_{WIL} ($\sigma_{\text{WIL}} = \sigma_{\text{MOI}} = 0$, $\delta_{\text{MOI}} = 0$).

As the results of Table B2 indicate, the performances of all the heuristics tend to degrade as δ_{WIL} increases. While the optimality ratio for the policies M, MW and MM degrades from approximately

3% when $\delta_{\text{WIL}} = 0$ to about 10% when $\delta_{\text{WIL}} = 30$, the performance of CE is quite robust and the error does not go above 2.6% - the effect which, as in the base scenario, is largely due to strong inertia of local supply corrections to bunkers prices.

The effect of the supply correction inertia η_{WIL} . As in the previous setting, we use $\sigma_{\text{WIL}} = \sigma_{\text{MOI}} = 0$ and $\delta_{\text{MOI}} = 0$. We now fix δ_{WIL} at its base-case value of 29.2 and analyze how the relative performance of the heuristics is influenced by changes in the value of η_{WIL} . As before, the prices at Moin are exclusively driven by the oil price P_0 while the prices at Wilmington are driven by both P_0 and the supply correction α_{WIL} .

η_{WIL}	$R^{\text{M}}(\%)$	$R^{\text{MW}}(\%)$	$R^{\text{MM}}(\%)$	$R^{\text{CE}}(\%)$
0.1	10.70	11.17	10.39	4.41
0.5	10.94	11.40	10.63	4.13
0.6	10.94	11.40	10.63	4.13
0.7	10.94	11.40	10.63	4.12
0.8	10.93	11.39	10.62	4.08
0.9	10.70	11.17	10.39	3.93

Table B3: Relative performance of heuristics as a function of η_{WIL} ($\delta_{\text{WIL}} = 29.2$, $\sigma_{\text{WIL}} = \sigma_{\text{MOI}} = 0$, $\delta_{\text{MOI}} = 0$).

In Table B3, we observe that changes in η_{WIL} have overall small effects on the performance of the four heuristics. However, it is interesting to note that the results presented confirm the intuition about the CE heuristic. As η_{WIL} increases, the supply correction α_{WIL} varies in a relatively slow fashion, which improves the performance of the CE heuristic. This is what we observe in the last column of Table B3.

The effect of volatility characterized by σ_{WIL} , σ_{MOI} , δ_{WIL} , δ_{MOI} . In Table B4, we assume $\delta_{\text{WIL}} = \delta_{\text{MOI}} = 0$, set $\sigma_{\text{WIL}} = \sigma_{\text{MOI}}$, and analyze the performance of the four policies as the value of σ_{WIL} is increased. Note that in this setting, the prices at Moin and Wilmington are driven by the oil price P_0 and a daily random correction in the absence of any supply corrections. On the other hand, in Table B5, we fix the values of η_{WIL} and η_{MOI} at their base-case levels of 0.99 and 0.97, respectively, set $y = \delta_{\text{WIL}} = \delta_{\text{MOI}} = \sigma_{\text{WIL}} = \sigma_{\text{MOI}}$ and analyze the performance of the heuristics as the value y is increased.

The results of Table B4 indicate that all heuristics have comparable performances. However, Table B5 identifies the CE heuristic as the best performer. This performance advantage stems from the fact that the CE heuristic is able to exploit high values of $\delta_{\text{WIL}} = \delta_{\text{MOI}}$ combined with high inertia of the local supply correction factors.

σ_{WIL}	$R^{\text{M}}(\%)$	$R^{\text{MW}}(\%)$	$R^{\text{MM}}(\%)$	$R^{\text{CE}}(\%)$
5	3.30	3.47	3.12	2.92
10	5.93	6.11	5.74	5.54
15	8.85	9.03	8.65	8.44
20	11.96	12.15	11.76	11.54

Table B4: Relative performance of heuristics as a function of $\sigma_{\text{WIL}} = \sigma_{\text{MOI}}$ ($\delta_{\text{WIL}} = \delta_{\text{MOI}} = 0$).

y	$R^{\text{M}}(\%)$	$R^{\text{MW}}(\%)$	$R^{\text{MM}}(\%)$	$R^{\text{CE}}(\%)$
5	4.10	4.26	3.91	2.56
10	7.66	7.84	7.47	4.80
15	11.60	11.79	11.40	7.34
20	15.88	16.09	15.68	10.09

Table B5: Relative performance of heuristics as a function of $y = \sigma_{\text{WIL}} = \sigma_{\text{MOI}} = \delta_{\text{WIL}} = \delta_{\text{MOI}}$ ($\eta_{\text{WIL}} = 0.99, \eta_{\text{MOI}} = 0.97$).