

University of Groningen

An inventory control policy for liquefied natural gas as a transportation fuel

Lopez Alvarez, Jose A.; Buijs, Paul; Kilic, Onur A.; Vis, Iris F. A.

Published in:
Omega: The International Journal of Management Science

DOI:
[10.1016/j.omega.2018.10.015](https://doi.org/10.1016/j.omega.2018.10.015)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):
Lopez Alvarez, J. A., Buijs, P., Kilic, O. A., & Vis, I. F. A. (2020). An inventory control policy for liquefied natural gas as a transportation fuel. *Omega: The International Journal of Management Science*, 90, Article 101985. <https://doi.org/10.1016/j.omega.2018.10.015>

Copyright


Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.



An inventory control policy for liquefied natural gas as a transportation fuel[☆]

Jose A. Lopez Alvarez*, Paul Buijs, Onur A. Kilic, Iris F.A. Vis

Department of Operations, Faculty of Economics and Business, University of Groningen, P.O. Box 800, Groningen 9700 AV, The Netherlands

ARTICLE INFO

Article history:

Received 14 February 2018

Accepted 22 October 2018

Available online 26 October 2018

Keywords:

Inventory control

Deteriorating inventory

Markov decision process

Liquefied natural gas

ABSTRACT

In this paper, we study a novel stochastic inventory management problem that arises in storage and refueling facilities for Liquefied Natural Gas (LNG) as a transportation fuel. In this inventory problem, the physio-chemical properties of LNG play a key role in the design of inventory policies. These properties are: (1) LNG suffers from both quantity decay and quality deterioration and (2) the quality of on-hand LNG can be upgraded by mixing it with higher-quality LNG. Given that LNG quality can be upgraded, an inventory control policy for this problem needs to consider the removal of LNG as a decision variable. We model and solve the problem by means of a Markov Decision Process (MDP) and study the structural characteristics of the optimal policy. The insights obtained in the analysis of the optimal policy are translated into a simple, though effective, inventory control policy in which actions (i.e., replenishment and/or removal) are driven by both the quality and the quantity of the inventories. We assess the performance of our policy by means of a numerical study and show that it performs close to optimal in many numerical instances. The main conclusion of our study is that it is important to take quality into consideration when design inventory control policies for LNG, and that the most effective way to cope with quality issues in an LNG inventory system involves both the removal and the replenishment of inventories.

© 2018 Elsevier Ltd. All rights reserved.

1. Introduction

Concerns about the environmental, economic and social effects of heavily polluting fossil fuels have driven governments around the globe to stimulate the development and use of alternative fuels. Among the currently available alternative fuels, Liquefied Natural Gas (LNG) is a prominent substitute for conventional fuels, especially in the maritime and long-haul road transportation sectors (Council Directive 2014/94/EU). Its main source, natural gas, is a low-cost energy source that is widely available and less polluting than fuel oils.

The adoption of LNG as a fuel in road and maritime transportation is strongly dependent on its competitiveness compared to conventional fuels in terms of cost, availability and fuel efficiency [9,20,25]. Therefore, a good design and efficient management of the LNG supply chain is paramount. Much of the extant research on LNG supply chain optimization has focused on the large-scale, global LNG supply chain, including applications in the maritime inventory routing problem for LNG [2,13,23]. Truck and ship operators using LNG as a fuel rely on the small-scale LNG sup-

ply chain, comprised of refueling stations and intermediary storage facilities. Research aimed at optimizing the small-scale LNG supply chain is scarce [10,16,20].

Supplying truck and ship operators with good quality LNG at a low cost is challenging due to the physio-chemical properties of LNG. To maintain its liquid state, LNG needs to stay cold. The boiling point for LNG is -162°C (-259°F) and from roughly that temperature LNG starts boiling off. The boil-off not only reduces the quantity of LNG in stock but also reduces its quality over time [17]. Trucks and ships refueling their tanks with LNG at a refueling station require a minimum fuel quality requirement. Below that requirement, their engines may run up to 5% less efficiently and, in extreme cases, can even be permanently damaged.

In this paper, we address the problem of maximizing the profit of an LNG storage or refueling facility by making replenishment and removal decisions for the LNG inventory. We study a single-location inventory system consisting of an LNG storage tank that must be replenished in order to meet uncertain customer demand for LNG. Unfulfilled demand is lost as no back-orders are allowed. The quality of the LNG stored in the tank must meet a known minimum quality requirement. This quality requirement is enforced by the market and all demand will be lost when the requirement is not met. An inventory control policy thus needs to consider the physio-chemical properties of LNG; that is the quantity decay and

[☆] This manuscript was processed by Associate Editor Kimms.

* Corresponding author.

E-mail address: j.a.lopez.alvarez@rug.nl (J.A. Lopez Alvarez).

quality deterioration that stem from the boil-off. Accordingly, if LNG is stored in the refuelling facility for an extended period of time, its quality might fall below the minimum quality requirement. An inventory control policy for this problem can benefit the fact that LNG is a liquid substance whose quality can be upgraded by mixing the LNG in stock with fresh LNG from a replenishment order. This implies that LNG needs not necessarily be removed from the inventory system when its quality drops below the quality requirement.

The problem described above can broadly be classified as an inventory control problem for commodities with quantity decay and quality deterioration, whose quality can be upgraded with replenishment orders. Prior research in the field of inventory problems for deteriorating products roughly falls into two groups: (1) considering quantity decay, i.e., where a fraction of the products is lost every period of time; and (2) considering quality deterioration, i.e., where products follow a quality deterioration process towards a state in which they are no longer marketable for their original purpose [19]. Most of the prior studies in this field consider either quality deterioration or quantity decay (for comprehensive surveys see, [4,12,15]). Since LNG suffers from both, our problem relates to the few studies that, in some form, have simultaneously considered quality deterioration and quantity decay.

Quality deterioration has been incorporated in inventory problems with quantity decay by, for example, specifying a maximum lifetime for products and using a time-varying quantity decay (e.g., [24,27,28]). In these studies, the rate of decay increases over time and becomes 100% when the maximum lifetime is reached. The notion of having maximum lifetime is relevant to our LNG inventory problem, as LNG has a minimum quality requirement, below which it cannot be sold. Yet, the quality deterioration process of LNG does not result in a (time-varying) quantity decay. Rather, quality deterioration occurs concurrently with a decay in quantity. Other studies have incorporated quality deterioration by modeling it as a decreasing value of the product over time [5–7,22]. Hence, in these problems, pricing decisions are an integral part of the replenishment decisions, while in the current fuel market, the price of LNG is only related to quality when it drops below the quality requirement. The notion of a minimum quality requirement is discussed in [18], modeling an inventory system for a decaying product where a penalty is incurred when sold below a certain weight. Minimum quality requirements also apply when selling LNG as a fuel. However, LNG below the quality requirement cannot be sold at all without upgrading.

We extend prior work by considering the option to upgrade the quality of inventories in stock by means of replenishment of inventories. Prior research on deteriorating inventory control considered a context in which the quality deterioration of the underlying goods was irreversible. The ability to upgrade the quality of LNG implies that, contrary to regular perishable commodities, the inventories do not need to be automatically removed from the system when their quality is below the minimum quality requirement accepted by the market. Accordingly, in the LNG inventory problem, the removal of inventories from the system is a decision.

In this paper, we model our problem as a Markov decision process (MDP). We use the MDP to obtain optimal policies for different numerical instances. By analyzing the resulting policies, we gain insights on optimal policies for the general problem; these insights will be illustrated using two examples. Based on these insights, we propose an inventory policy for our problem. The performance of the policy is evaluated in various numerical instances.

2. Problem description

We study the inventory problem faced by a refueling or storage facility in the small-scale LNG supply chain. This inventory sys-

Table 1
Input parameters.

M	Storage tank size
D	Stochastic demand per period
W_-	Minimum quality requirement
W_+	Supplier's LNG quality
b	Quality deterioration per period
θ	Quantity decay per period
r	Revenue per unit of LNG sold
K	Fixed ordering cost
c	Cost per unit of LNG
g	Cost/revenue for each unit of LNG removed
h	Holding cost per unit
ρ	Discount factor

tem consists of a single stock point that is periodically reviewed. This stock point is an insulated LNG storage tank with a maximum capacity of M units of volume. Every period, the facility faces a random demand D , which follows a known probability distribution with mean μ and standard deviation σ . Due to the characteristics of LNG-fueled engines, demand for LNG can only be fulfilled if the LNG held in the storage tank complies with the minimum quality requirement W_- established by the market. If this quality requirement is met, the LNG held in the storage tank will be used to fulfill the demand; otherwise, all the demand of the period will be lost; i.e., no back-orders are allowed.

The LNG quality level is measured as an absolute value representing the methane number of the fuel. The methane number is a well-established quality measure for the use of LNG as a fuel. In order to comply with the minimum quality requirement, inventory management decisions need to consider that LNG boils off. Boil-off is a natural process in which the LNG vaporizes due to the ambient heat input. This process induces quantity decay and quality deterioration of the LNG in stock. The daily quantity decay and quality deterioration of the fuel depend on several complex physio-chemical processes [17], which can be reasonably approximated by representing quality deterioration as a fixed volume b per period and quantity decay as a fixed amount θ per period. For more in-depth information about the physio-chemical properties of LNG and our assumptions in that regard, we refer to Appendix A.

We consider a setting where a single LNG supplier offers LNG with fixed and known quality W_+ , where $W_+ \geq W_-$. Every time the facility is replenished, the LNG held in the storage tank – if any – is mixed with the incoming LNG from the replenishment order. The mixture of both loads not only increases the inventory level but also changes the quality of the on-hand LNG. We model the resulting quality of the mixture using the weighted average of the quality of the loads involved in the mixture (see Appendix A for a justification of this assumption). This implies that the quality of the LNG held in the storage tank can be upgraded by means of a replenishment order. In fact, this property can be used to upgrade the quality of LNG whose quality is below the minimum quality requirement. Accordingly, the fuel need not necessarily be removed from the inventory system when its quality drops below the minimum quality requirement; instead, the removal of inventories from the system is a decision variable.

Taking both the inventory level and the quality of LNG into consideration, two decisions are to be made: (1) to determine the amount of LNG to remove from the storage tank and (2) to decide the order quantity. Motivated by the fact that replenishment of the facility can be done overnight, we assume there is no lead time for replenishment. Additionally, we assume that the removal of LNG can be done instantaneously and that it occurs before replenishment takes place. Removing the LNG before replenishment is advantageous from a practical point of view because the quality of the LNG stored in the tank is lower before it is mixed

with the “fresh” LNG coming from replenishment. The order of events/activities throughout a period in the system is as follows:

1. The state of the system (i.e., the inventory level and the quality of the LNG) is reviewed
2. Based on the system state the planner decides:
 - (a) The amount of LNG to remove from the system
 - (b) The order quantity
3. The removal of LNG takes place, according to step 2a
4. Arrival of the replenishment order, according to step 2b. The LNG coming from the replenishment order is mixed with the remaining LNG in the storage
5. The demand of the period is realized
6. The quantity decay and the quality deterioration of the period take place

The objective in this problem is to determine an inventory policy that maximizes the expected discounted profit of the refuelling facility for an infinite planning horizon. That is the sum of the expected profit of the facility for the infinite horizon, where each future profit is discounted according to a discount factor ρ . For every volume unit of LNG sold, the facility receives a fixed revenue r . There is a fixed replenishment cost K each time a replenishment order is placed and a variable cost c per unit volume of LNG purchased. Every unit of LNG removed from the inventory system generates a cost of g . Since some refuelling stations can make revenue with the removed LNG, the value of g can be negative as long as it does not exceed the unit cost ($-g < c$). Each volume unit of LNG carried after the demand of the period takes place incurs in a holding cost h .

3. Model formulation

In order to obtain an optimal inventory policy, we model the problem as an infinite horizon discrete-time Markov decision process. The MDP model is characterized by its states, actions, transition from one state to another, profit and value function. We describe this model below.

States: We define the state of the system at the beginning of period t as $Z_t = (I_t, W_t)$, where I_t is the inventory level and W_t the quality of the on-hand inventory. The inventory volume is bounded by the tank size ($0 \leq I_t \leq M$), and the quality of the inventory W_t cannot exceed the quality provided by the supplier ($0 \leq W_t \leq W_+$).

Actions: Upon reviewing the state of the system at the beginning of period t , decisions need to be made on the amount Y of LNG to be removed and the amount X of LNG to be replenished. The amount of LNG that can be removed is at most the amount of LNG stored in the facility ($0 \leq Y \leq I_t$), and the replenishment order size cannot exceed the remaining space available in the storage tank after removal ($0 \leq X \leq M - (I_t - Y)$).

State transitions: The state transition between two subsequent periods t and $t + 1$, can be described in two phases. In the first phase, the system state is updated based on the actions, which take place at the beginning of period t . Specifically, the first phase covers steps 1 through 4 of the daily operations (as specified in Section 2). We define $Z'_t = (I'_t, W'_t)$ as the post-action state that describes the status of the system immediately after the first phase is completed. The transition $Z_t = (I_t, W_t)$ to $Z'_t = (I'_t, W'_t)$ is as follows:

$$Z'_t = \left(I_t - Y + X, \frac{(I_t - Y) \cdot W_t + (X \cdot W_+)}{I_t - Y + X} \right). \quad (1)$$

In this expression, the inventory level is updated based on the quantity of LNG removed and replenished. The update in quality is determined by the weighted average of the quality of the LNG in stock and the quality of the LNG coming from the replenishment order. Note that the resulting quality of the mixture does not

include the LNG that was removed from the system since removal precedes replenishment.

In the second phase, the system state is updated based on demand, quantity decay, and quality deterioration (steps 5 and 6 as specified in Section 2). The demand and boil-off are subtracted from the inventory level if the quality is higher than the minimum quality requirement; if that requirement is not met, all demand will be lost and the system state will only be updated based on the quantity decay and quality deterioration. Let D_t be the random demand in period t . The transition $Z'_t = (I'_t, W'_t)$ to $Z_{t+1} = (I_{t+1}, W_{t+1})$ is as follows:

$$Z_{t+1} = \begin{cases} (\max\{I'_t - D_t - \theta, 0\}, \max\{W'_t - b, 0\}), & \text{if } W'_t \geq W_- \\ (\max\{I'_t - \theta, 0\}, \max\{W'_t - b, 0\}), & \text{if } W'_t < W_- \end{cases} \quad (2)$$

In this expression, the inventory level I'_t is updated based on the demand (if $W_t \geq W_-$) and the quantity decay θ . The quality update is entirely based on the deterioration rate b . The reason why the max operators are included in the expression is that they bound the inventory level and the quality, so they do not fall below zero.

Profit function: The daily profit of the system is composed of revenues, ordering costs, purchasing costs, removal costs (revenues) and holding costs. We define $\beta(X, Y)$ as the sum of the costs (revenues) induced by the actions of the system.

$$\beta(X, Y) = \begin{cases} K + (c \cdot X) + (g \cdot Y) & \text{if } X > 0, \\ g \cdot Y & \text{if } X = 0. \end{cases} \quad (3)$$

Furthermore, we define $F(S'_t)$ as the function that represents the expected amount of LNG sold in a period and $H_t(S'_t)$ as the expected inventory level after demand takes place.

$$F(Z'_t) = \begin{cases} E[\min(D_t, I'_t)] & \text{if } W'_t \geq W_- \\ 0, & \text{if } W'_t < W_- \end{cases} \quad (4)$$

$$H(Z'_t) = I'_t - F(Z'_t). \quad (5)$$

If $W'_t < W_-$, the LNG cannot be sold and, therefore, the function takes the value zero. We define the expected profit of the system in period t as:

$$P(Z_t, X, Y) = r \cdot F(Z'_t) - h \cdot H(Z'_t) - \beta(X, Y). \quad (6)$$

The profit of the system is the sum of the revenues obtained from the LNG sold, minus the holding costs and the costs(revenues) induced by the actions. The holding costs are computed based on the on-hand inventory after the demand takes place.

Value function: The function $V(S_t)$ represents the optimal discounted profit of the system, which is the function to be optimized in our model. We define this function as:

$$V(Z_t) = \max_{(X,Y)} \{P(Z_t, X, Y) + (1 - \rho) \cdot E[V(Z_{t+1})]\}. \quad (7)$$

This function is composed of the costs of the immediate period plus the discounted expected value of the states reached in the following period.

3.1. MDP implementation

In order to obtain the optimal removal and replenishment policy, we solve the MDP using the value iteration algorithm [21]. Due to the discrete nature of this solution approach, we redefine the states and the actions of our problem as discrete sets:

$$Z = \{(I, W) \mid I=0, \Delta_I, 2\Delta_I, \dots, M \text{ and } W=0, \Delta_W, 2\Delta_W, \dots, W_-\}, \quad (8)$$




Fig. 1. Classification of states based on optimal actions for Example 1.

and,

$$A_{\{Z_t\}} = \{(X, Y) \mid Y = 0, \Delta_I, 2\Delta_I, \dots, I_t - Y \text{ and } X = 0, \Delta_I, 2\Delta_I, \dots, M - (I_t - Y)\}, \quad (9)$$

where Δ_I and Δ_W are the interval sizes of the inventory levels and the quality respectively.

Throughout our numerical study, we assume that we can model demand with high precision by using a discrete random variable. Furthermore, taking into consideration that the weighted average is used to compute the quality of a mixture of LNG loads, it is possible that the resulting quality does not belong to the discrete scale specified above; if that is the case, we round down the resulting quality to a quality value specified in the discrete scale. Finally, throughout the numerical analysis, we assume that in the initial condition of the system the inventory level is zero.

4. Illustrative example of an optimal policy

In this section, we explore how optimal policies of our MDP behave by means of two illustrative examples. In Example 1, we consider a numerical instance where no revenue is obtained when LNG is removed from the system (i.e., $g = 0$), while in Example 2, we consider the case where the removal of LNG generates a revenue (i.e., $g = -20$). All other input parameters are equal in both examples and are the following: the maximum capacity of the storage tank is $M = 50$ and the demand per period follows a negative binomial distribution with mean $u = 8$ and standard deviation $\sigma = 4$ units of volume. Replenishment orders are placed to a supplier whose LNG has a quality of $W_+ = 24$. The fixed cost of placing a replenishment order is $K = 200$ and the cost per unit of volume is $c = 35$. In order to fulfill demand, the LNG of the refueling station needs to comply with the minimum quality requirement, which is $W_- = 8$. Furthermore, the daily quality deterioration of the LNG is $b = 8$. We use the following values for the other parameters: $\theta = 1$, $p = 130$, $\Delta_I = 1$, $\Delta_W = 1$, $g = 0$, $h = 1$ and $\rho = 0.05$.

Analysis of the states in which a replenishment and/or removal action is needed. Figs. 1 and 2 are a graphical representation of the states in which actions need to be carried out under an optimal policy for each of the illustrative examples. In these figures, states in which no action is required are represented by dots and

states in which an action is required are represented by either circles, triangles or diamonds. States represented with a diamond are those in which no removal takes place, but a replenishment order is placed ($Y = 0$ and $X > 0$); states represented with a circle are those in which both removal and replenishment are performed ($Y > 0$ and $X > 0$); states represented with a triangle are those in which only removal takes place ($Y > 0$ and $X = 0$). The filling color of each symbol shows whether the probability of reaching that state under the optimal policy is larger than zero (in black) or zero (in grey). Note that this probability only considers the likeliness of reaching a state at the end of a period; hence, post-action states are not considered in this probability. Furthermore, we truncated both the quality and the inventory axes, at levels 20 and 35 respectively, because the states on those truncated areas are unreachable under the optimal policy of each of the examples.

We observe that the policy for both examples has a tendency to have two threshold values that trigger actions in the system; one quantity-based and one quality-based, which we define as \bar{I} and \bar{W} , respectively. Both these thresholds are depicted in Figs. 1 and 2. When either the quality is lower than \bar{W} or the on-hand inventory level is lower than or equal to \bar{I} , the policy tends to take a replenishment and/or removal action. In both examples, the value of the quality threshold \bar{W} is equal to 8, while the quantity threshold \bar{I} is equal to 6 in the first example and 5 on the second example. A reasonable interpretation of these thresholds is that \bar{I} helps to prevent the system from losing demand due to lack of inventory, while threshold \bar{W} helps to prevent low-quality inventories. In fact, we observe that $\bar{W} = W_-$, which implies that an optimal policy responds immediately after the quality falls below the minimum quality requirement. Hence, we observe that under an optimal policy, W is never below $W_- - b$ due to the immediate response of the policy to quality issues.

We can also see that the states represented with a triangle in both examples are not driven by the aforementioned threshold levels. We interpret the removal action in these states as a means to reduce unnecessary holding costs for LNG that most likely will have to be removed once the threshold level \bar{W} is exceeded. Accordingly, this type of states is more likely to appear in numerical instances where the removal of LNG generates a revenue, which explains why Example 2 has more of this type of states than Example 1. Conversely, in those instances in which the removed LNG




Fig. 2. Classification of states based on optimal actions for Example 2.

Table 2
Inventory before and after removal of LNG ($W = 4$) for both examples.

State $Z = (I, W)$		Example 1		Example 2	
I	W	Removal (Y)	Inv. post removal ($I - Y$)	Removal (Y)	Inv. post removal ($I - Y$)
14	4	0	14	5	9
15	4	0	15	6	9
16	4	0	16	7	9
17	4	1	16	8	9
18	4	2	16	9	9

Table 3
Inventory before and after removal of LNG ($W = 0$) for both examples.

State $Z = (I, W)$		Example 1		Example 2	
I	W	Removal (Y)	Inv. post removal ($I - Y$)	Removal (Y)	Inv. post removal ($I - Y$)
12	0	0	12	5	7
13	0	1	12	6	7
14	0	2	12	7	7
15	0	3	12	8	7
16	0	4	12	9	7
17	0	5	12	10	7

generates a cost, the presence of this type of states in the optimal policy is less common.

Analysis of the removal action. Removal of LNG mostly takes place in some of the states where the quality of the LNG has fallen below threshold \bar{W} . We observe that for the majority of these states, the removal action generally reduces the volume of LNG to a fixed inventory level that seems to depend on the current quality level of the LNG. We interpret this behavior as a “remove down to level”. Tables 2 and 3 present the removal decision for a number of states in Examples 1 and 2.

In Table 2 we see that, for a given LNG quality level of $W = 4$ in Example 1, removal mostly takes place if the inventory level is higher than 16. If so, the inventory immediately after removal ($I - Y$) is equal to 16. A similar pattern can be observed in Table 3, where $W = 0$; however, the removal down to level is 12 in that case. Further analysis of the other states in the illustrative examples suggests that the removal down to level is dependent on the quality. It seems that the lower the quality, the lower the value of the removal down to level. Accordingly, lower quality tends to result in greater removal of LNG from the system.

It is important to remark that not all removal decisions that occur when threshold \bar{W} is exceeded are driven by the remove down to level. In both examples, the optimal policy dictates that all the LNG is removed from the system and that a replenishment order must be placed when the inventory level is low (i.e., 1 unit for Example 1 and 3 or fewer units for Example 2). The removal of this low amount of LNG is done to ensure that the quality of the LNG upon replenishment is equal to W_+ . If that LNG was not removed, the quality after replenishment would be strictly lower than W_+ .

The analysis of both illustrative examples suggests three potential reasons why removal is part of an optimal policy. The first reason is that the removal of LNG reduces the amount of low-quality LNG that will be mixed with the LNG that comes from the replenishment order. This enhances the quality of the resulting LNG mixture. The second reason is that the removal of LNG could prevent a temporal dead-lock situation in which the facility will neither be able to serve demand (due to the low quality of the LNG inventories) nor upgrade the quality of the LNG to meet the minimum quality requirement (due to a storage tank that is already filled close to maximum capacity). Note that this temporal dead-

Table 4
States, actions and post-action states of Example 1.

State $Z = (I, W)$		Removal	Replenishment	Post-action state $Z' = (I', W')$	
I	W	Y	X	I'	W'
0	0	0	27	27	24
1	0	1	27	27	24
2	0	0	19	21	21
3	0	0	18	21	20
16	4	0	4	20	8
17	4	1	4	20	8
18	4	2	4	20	8
19	4	3	4	20	8
1	10	1	27	27	24
2	10	0	19	21	22
3	10	0	18	21	22

Table 5
States, actions and post-action states of Example 2.

State $Z = (I, W)$		Removal	Replenishment	Post-action state $Z' = (I', W')$	
I	W	Y	X	I'	W'
1	0	1	26	26	24
2	0	2	26	26	24
3	0	3	26	26	24
4	0	0	16	20	19
15	4	6	14	23	16
16	4	7	14	23	16
17	4	8	14	23	16
1	10	1	26	26	24
2	10	2	26	26	24
3	10	3	26	26	24
4	10	0	16	20	21

lock may resolve itself eventually due to the quantity decay (unless the quantity decay θ is zero). The third reason is that there might be situations in which there is an excess of LNG inventories which are expected to be removed eventually. Hence, the removal of some of the on-hand LNG would reduce the holding costs.

Analysis of the replenishment order size. We observe in both examples that the size of the replenishment order seems to be affected by both the inventory level and the quality of the on-hand inventories. It seems that replenishment decisions are not driven by either a fixed order quantity, an order up to level, or some sort of quality up to level (i.e., referring to a fixed post-action quality level). The quantity and quality of the LNG after replenishment differ significantly depending on the state of the system before replenishment. To better illustrate this, Tables 4 and 5 present the optimal decisions and the post-action state for a number of states in Examples 1 and 2 respectively. When comparing two states with the same quality $W = 0$ and different inventory $I = 0$ and $I = 3$ in Table 4, we see that the replenishment order size, the post-action inventory level and the post-action quality are higher for the case where $I = 0$. Although there is no a clear tendency that defines the replenishment decision, it is interesting to see that when $I - Y \neq 0$, the inventory level post-action seem to be quite similar regardless of the quality. This might suggest that in states when $I - Y > 0$, a single order up to level might not be too far from an optimal decision.

Taking into consideration insights from the illustrative examples, it appears that:

1. The majority of actions in the system are triggered by a quantity and a quality threshold level
2. Most of the removal actions take place when the quality of the on-hand LNG is below the quality threshold level
3. When removal actions are required, those decisions are usually driven by a remove down to level, which differs for specific quality levels

4. The size of replenishment orders does not follow a simple decision rule

5. Special cases

Throughout our analysis, we identified a few special cases in which the optimal policy for our problem can be obtained using solution approaches for existing problems in the literature.

1. *Case $W_+ = W_-$.* In this special case, the quality of the supplier's LNG is equal to the minimum quality requirement. This implies that the quality of the LNG in stock at the end of a period will always be below the minimum requirement. Since $W_+ = W_-$, it is not possible to upgrade the quality of the LNG in stock to a level above or equal the minimum requirement. Consequently, all the remaining LNG needs to be removed from the system. The resulting problem resembles a news vendor problem for a decaying commodity with fixed ordering cost and holding cost. To solve this problem, we can write the profit function as:

$$f(X) = r \cdot E[\min(D, X)] - (h + g) \cdot E[\max(X - D - \theta, 0)] - c \cdot X - K. \quad (10)$$

Similarly to the cost function of a standard news vendor problem (see [3]), the function above is concave. Accordingly, the optimal replenishment order size can be obtained by identifying the critical point of the profit function.

2. *Case $b = 0$ or $W_- = 0$.* In this case, the quality of the LNG never falls below the minimum requirement. The resulting problem can be modeled as a periodic inventory problem with stochastic demand, lost sales and quantity decay. In this problem, instead of maximizing the expected discounted profit, one should minimize expected discounted cost. It was shown in [26] that (S, s) policies are optimal for this problem when the quantity decay is a fraction of the remaining inventory at the end of each period. Since the quantity decay in our problem is a constant amount of LNG, the properties of the optimal cost function are not affected. Accordingly, an (S, s) policy is optimal for this special case of our problem.
3. *Case with deterministic demand.* When demand is considered to be deterministic, the solution to the LNG inventory problem is trivial. In this setting, an optimal policy would not require any safety stock to buffer against stochastic demand. Accordingly, under the optimal solution, there would not be mixing of LNG, which means that LNG can be considered as a regular perishable commodity with quantity decay. The solution to this problem is a simple policy as shown in [19].

6. Definition of an inventory policy

In this section, we introduce the (S, s, v, k) policy, which is an inventory policy inspired by the insights obtained from analyzing an optimal policy. The (S, s) policy is used as a starting point to develop our policy. The reason is that the (S, s) policy is (1) a policy that is easily implementable in practice, (2) optimal in certain special cases of the LNG inventory problem as shown in Section 5 and (3) captures the quantity threshold of the optimal policy with parameter s .

We extend the (S, s) policy by introducing parameter v . This parameter represents a quality threshold level, which serves as a trigger to respond to situations in which the LNG cannot be sold because its quality is below the minimum requirement. We learned in the illustrative example of Section 4 that the response driven by the quality threshold v is to place a replenishment order and, sometimes, remove some LNG. The removal decision in the (S, s, v, k) policy is determined by parameter k . This parameter serves as a removal down to level, which is a control rule used in other

inventory problems that incorporate removal decisions (e.g., [14]). Specifically, when the quality level W is below v and the inventory level I is higher than k , then $I - k$ units of LNG need to be removed from the system. Regardless if removal is needed or not, a replenishment order needs to be placed when threshold v is exceeded, such that the post-action inventory level is equal to S .

We analyzed various numerical instances in order to identify any properties that can be used to obtain the optimal parameters of the (S, s, v, k) policy. First of all, in all numerical examples considered we observed that the value of parameter v is always equal to the minimum quality requirement under an optimal policy. This, however, is not necessarily the case when the other parameters $(S, s$ and $k)$ are not optimal. For example, when the value of s is forced to be zero, then the parameter k can take a value higher than the minimum requirement to compensate for the lack of a quantity threshold trigger. Furthermore, we studied the behavior of the profit function of our problem with respect to the parameter k . We noticed that the profit function can have several critical points when fixing the values of the S, s and v . Therefore, in order to compute optimal parameters of the (S, s, v, k) policy throughout our numerical experiments, we make a full enumeration of all four parameters.

7. Numerical study

The purpose of the numerical study is twofold: (1) to evaluate and analyze the performance of the (S, s, v, k) policy and (2) to test the sensitivity of the optimal expected profit of the system with respect to the input parameters of the problem.

For our numerical analysis, we introduce parameter ϕ , which is an indicator of quality criticality. Specifically, ϕ represents the number of periods it takes for a volume unit of LNG with quality W_+ to fall below the minimum quality requirement W_- , i.e., the lower the value of ϕ , the more critical LNG quality is in our problem:

$$\phi = \frac{W_+ - W_-}{b} + 1. \tag{11}$$

7.1. Performance of the (S, s, v, k) policy

In this section, we test the performance of the (S, s, v, k) policy and two other simple policies against the optimal policy. The policies to be considered are the following:

1. An optimal policy obtained from the MDP.
2. An (S, s, v, k) policy
3. An $(S, s, v, k = M)$ policy
4. An $(S, s, v, k = 0)$ policy

Policies 3 and 4 represent a specific group of (S, s, v, k) policies in which the removal decision is simplified. In policy 3 we force parameter k to be equal to the storage tank size M , which means that in this policy no removal of LNG takes place. This policy might be useful in practice for LNG facilities in which removal of LNG is not possible or economically intractable. Furthermore, by studying the performance of an inventory policy that does not incorporate removal decisions, we can gain insights into the relevance of the removal of LNG for an inventory policy for the problem at hand. In policy 4, we enforce that all inventories are removed when quality falls below the threshold v (hence, $k = 0$). In this policy, the LNG is thus treated as a regular perishable commodity in the sense that it is automatically removed from the inventory system as soon as its quality falls below the minimum quality requirement.

Experimental design. The performance of the policies is tested in various experiments considering a range of values for the input parameters of our problem as shown in Table 7. The complete set

Table 6
Numerical design.

Parameter	Base value	Min value	Max value	Step increment
ϕ	—	1	5	0.5
b	4	4	16	4
M	30	15	45	5
θ	0.5	0.5	1	0.5
μ	8	8	14	2
σ	4	4	6	0.5
K	200	150	300	50
c	40	40	70	10
h	1	1	5	1
p	120	120	180	20
g	0	-20	20	20

of experiments was created by changing the values of one parameter at a time, whilst the other parameters were kept on their base value. The only exception is ϕ ; we tested all the values of this parameter for all the instances considered. We opted to make ϕ a pivotal parameter in this experimental design because a preliminary analysis showed that this parameter has a strong effect on the performance of the policies. Throughout the experiments we fixed the value of $\rho = 0.1$, $\Delta_I = 0.25$, $\Delta_W = 0.25$. Furthermore, in all experiments the value of W_+ is equal to $W_- + b \cdot (\phi - 1)$ and the value of W_- is equal to b . Finally, we model demand using a negative binomial distribution with mean μ and standard deviation σ .

For each instance, we compute the optimal parameters of all the three types of (S, s, v, k) policies by fully enumerating the solution space. Furthermore, we measure the relative gap of each of the three types of (S, s, v, k) policies with respect to the optimal as follows:

$$\text{Relative gap} = \frac{\pi^{\text{opt}} - \pi^{\text{heu}}}{\pi^{\text{opt}}} \cdot 100, \tag{12}$$

where π^{opt} is the expected profit of the optimal policy and π^{heu} is the profit of the heuristic policy.

Results. Fig. 3 shows the average expected discounted profit of the system (henceforth referred as to profit) of the four policies for different levels of ϕ . These results show that the profit of all policies increases with ϕ , up to the point where it becomes insensitive to further increments of ϕ . This can be explained by the fact that the policies are less affected by quality considerations as the quality criticality decreases, up to the point where the inventory policy is purely driven by inventory levels.

It is noteworthy that two of the special cases presented in Section 5 are represented in Fig. 3. Firstly, when the quality criticality is low (i.e., the value of ϕ is high), the optimal policy is purely driven by inventory levels, which resembles the special case, where $b = 0$. As shown in Section 5, an optimal policy for this special case is an (S, s) policy, which explains why the profit of all four policies tends to converge when ϕ is large. Secondly, when $\phi = 1$, our problem is equivalent to the special case where $W_+ = W_-$. In fact, when $\phi = 1$, the profit of the system for both (S, s, v, k) and $(S, s, v, k = 0)$ is optimal. This is intuitive given the optimal policy for this special case entails removing all LNG from the system ($k = 0$) and placing a fixed replenishment order of size S - as in the news vendor problem. The $(S, s, v, k = M)$ policy performs poorly in this special case because it requires selecting a conservative order up to level to prevent a temporal dead-lock situation.

Table 7 shows the average gap between the profit of each of the three types of (S, s, v, k) policies and the profit of the optimal policy, for different values of ϕ . The table also shows the standard deviation of the gap and the maximum gap. The average gap of the (S, s, v, k) policy is 1.65%, while the gap of the $(S, s, v, k = 0)$ is 5.57% and that of the $(S, s, v, k = M)$ is 24.56%. Given that the




Fig. 3. Average expected profit for the different policies.

Table 7
Statistics of the performance of the policies with respect to the optimal.

ϕ	(S, s, v, k)			$(S, s, v, k = 0)$			$(S, s, v, k = M)$		
	Mean %	Std.Dev %	Max %	Mean %	Std.Dev %	Max %	Mean %	Std.Dev %	Max %
1.0	0.00	0.00	0.00	0.00	0.00	0.00	95.72	8.69	100.00
1.5	4.34	2.93	17.26	25.35	11.5	73.37	72.79	16.90	100.00
2.0	4.56	1.59	9.91	4.56	1.59	9.91	28.19	7.36	50.14
2.5	1.94	0.96	5.48	7.54	2.40	15.82	13.35	5.26	31.98
3.0	1.99	0.88	4.17	4.74	1.23	7.68	6.59	3.32	17.54
3.5	1.18	0.64	2.85	3.53	1.16	6.05	2.76	1.95	9.68
4.0	0.56	0.40	1.62	2.27	0.95	4.23	1.10	1.09	5.16
4.5	0.21	0.19	0.78	1.46	0.79	3.61	0.40	0.53	2.55
5.0	0.05	0.07	0.32	0.72	0.50	2.18	0.10	0.21	1.02
Average	1.65	—	—	5.57	—	—	24.56	—	—

$(S, s, v, k = 0)$ and the $(S, s, v, k = M)$ policies are special cases of the (S, s, v, k) policy, the performance of the latter is the highest among the three. For all policies, the performance with respect to the optimal tends to improve as the value of ϕ increases, up to the point where all policies tend to perform optimally. For example, the standard deviation of the (S, s, v, k) policy is 1.59% and 0.07% when $\phi = 2$ and $\phi = 5$ respectively. The exception to this pattern is the special case where $\phi = 1$, when both (S, s, v, k) and $(S, s, v, k = 0)$ are optimal. The $(S, s, v, k = M)$ policy performs particularly bad when $\phi < 2$ due to the lack of removal. In fact, in many of the instances where $\phi < 2$, the optimal order up to level S for the $(S, s, v, k = M)$ policy is zero, which explains why the maximum gap of that policy when $\phi < 2$ is 100%.


Analysis of the cost drivers. Throughout our numerical study we identified the following cost drivers: (a) the average amount of LNG purchased per time unit; (b) the average amount of LNG that is removed from the system; (c) the average (post-demand) inventory of the system and (d) the average replenishment frequency (i.e., average replenishment orders per period). In Fig. 4 we illustrate each of the cost drivers for different values of ϕ for all policies.

Fig. 4a shows that the amount of LNG purchased of all four policies. Here, we can observe that in the $(S, s, v, k = M)$ policy the amount of LNG purchased is much lower than that of the other policies. This is caused by two main reasons: (1) the inability


to remove LNG from the system forces the policy to purchase small quantities of LNG in order to increase the probability that all the inventories are sold before the LNG quality falls below the minimum requirement and (2) when $\phi < 2$, there are instances in which the amount of LNG purchased is zero.

When comparing the $(S, s, v, k = 0)$, the (S, s, v, k) and the optimal policies, we see that there are no major differences in the inventory levels as shown in Fig. 4c. All these policies sustain their inventory level with a similar amount of LNG purchased and also a similar replenishment frequency as shown in Figs. 4a and 4d respectively. The major difference among these policies relates to the amount of LNG removed as the $(S, s, v, k = 0)$ and (S, s, v, k) policies induce more removal than an optimal policy as shown in Fig. 4b. The reasoning behind the relatively low removal of LNG in an optimal policy is that this policy can tackle quality issues more effectively than the other policies due to its flexibility to select the amount of LNG to replenish and remove in different states.


To further analyze the difference in the amount of LNG that is removed among the policies, Fig. 5 presents the percentage of times in which a replenishment order was placed because the LNG fell below the minimum quality requirement W_{-} . In this figure, we observe that in the optimal policy quality falls below the minimum requirement more times as compared to the other policies. This is intuitive given that all three types of (S, s, v, k) policies are




(a) Average amount of LNG purchased



(b) Average amount of LNG removed



(c) Average inventory level



(d) Average Replenishment frequency

Fig. 4. Statistics of the performance of the policies.

more penalized than the optimal policy when quality falls below W_- . Specifically, in the case of the $(S, s, v, k = M)$, if the quality falls below W_- , it might induce an excess of low-quality inventory since this policy does not incorporate any removal of LNG. In the case of the other two policies, when quality falls below W_- , an excess of removal of inventory occurs (as shown in Fig. 4b), which is economically penalized by the removal cost (if any) and by the loss of the capital invested in those removed inventories. Interestingly, an optimal policy induces less removal than those policies, even though the quality of the on-hand LNG falls below the minimum quality requirement more times than in the other policies; this shows the ability of an optimal policy to cope with quality issues.

Finally, we observe that the behavior of the performance measures in Figs. 4 and 5 is not monotonic across different values of ϕ . This non-monotonic behavior is mostly explained by the periodicity of the problem. Since we consider a periodic inventory management problem, LNG quality basically becomes a shelf life, i.e., a number of periods the LNG remains above the minimum quality requirement. If that shelf life is, for example 1.8 periods, the LNG would effectively remain above the minimum quality requirement for one period. It is important to note that at integer values of ϕ ,

any on-hand inventory would immediately reduce the shelf life of the LNG mixture to a value strictly lower than ϕ . Therefore, the policies tend to avoid mixing LNG when ϕ has an integer value, especially for lower values of ϕ . At higher values of ϕ quality issues are less urgent, which may make it economically attractive to control quality by upgrading the quality of LNG by means of mixture, even at integer values of ϕ .

Analysis of the cost/revenue components. In Tables 8–10 we show the cost/revenue components for all policies for all values of ϕ . Since g can be revenue or a cost, we opted to present one table for each value of g considered our experimental design (i.e., $g = -20, 0, 20$). The numbers presented in each table represent an average over the numerical instances. In all tables, H represents the total inventory costs, P represents the purchasing costs, O represents the ordering costs, E represents the removal costs/revenue and R represents the revenue obtained from selling LNG. In general, we observe that the major cost component of all policies is the purchasing cost, while the holding cost is the smallest portion of the overall costs. On average, the holding costs represent 1.75% of the total cost of the policies, while the ordering costs represent 23.8% and the purchasing cost 74.4%. These proportions change to some extent along different values of ϕ . For instance, as the value




Fig. 5. Percentage of times in which replenishment orders are triggered by low-quality LNG.

Table 8

Cost and revenue components (in units of 1000 euros) for all policies when $g = 0$.

ϕ	Optimal					(S, s, v, k)					(S, s, v, k = 0)					(S, s, v, k = M)				
	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R
1.0	0.03	2.04	3.90	0.00	8.96	0.03	2.04	3.90	0.00	8.96	0.03	2.04	3.90	0.00	8.96	0.00	0.54	0.36	0.00	1.17
1.5	0.05	2.04	3.82	0.00	9.73	0.05	2.04	3.91	0.00	9.68	0.03	2.04	3.90	0.00	8.96	0.01	1.68	1.57	0.00	4.55
2.0	0.08	1.40	3.79	0.00	9.78	0.07	1.15	3.92	0.00	9.45	0.08	1.11	3.93	0.00	9.43	0.03	1.56	2.73	0.00	7.62
2.5	0.09	1.27	3.86	0.00	10.02	0.09	1.18	3.94	0.00	9.92	0.07	1.26	3.93	0.00	9.73	0.05	1.46	3.30	0.00	9.00
3.0	0.12	1.04	3.92	0.00	10.08	0.10	1.09	3.87	0.00	9.97	0.10	0.95	3.97	0.00	9.80	0.08	1.21	3.58	0.00	9.56
3.5	0.13	0.94	3.98	0.00	10.14	0.11	1.00	3.90	0.00	10.05	0.11	0.95	3.95	0.00	9.92	0.10	1.03	3.78	0.00	9.87
4.0	0.15	0.86	4.04	0.00	10.18	0.14	0.89	4.01	0.00	10.14	0.12	0.90	4.00	0.00	10.06	0.13	0.93	3.93	0.00	10.08
4.5	0.15	0.82	4.06	0.00	10.19	0.15	0.84	4.03	0.00	10.17	0.13	0.86	3.99	0.00	10.07	0.14	0.85	4.00	0.00	10.14
5.0	0.16	0.80	4.08	0.00	10.20	0.15	0.80	4.07	0.00	10.20	0.14	0.83	4.05	0.00	10.16	0.15	0.80	4.06	0.00	10.19

Table 9

Cost and revenue components (in units of 1000 euros) for all policies when $g = -20$.

ϕ	Optimal					(S, s, v, k)					(S, s, v, k = 0)					(S, s, v, k = M)				
	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R
1.0	0.03	2.00	4.20	0.52	8.70	0.03	2.00	4.20	0.52	8.70	0.03	2.00	4.20	0.52	8.70	0.00	0.00	0.00	0.00	0.00
1.5	0.04	2.00	3.75	0.19	9.10	0.04	2.00	3.81	0.23	9.05	0.03	2.00	4.20	0.52	8.70	0.01	2.00	1.65	0.00	4.63
2.0	0.06	1.22	3.80	0.18	9.03	0.07	1.09	4.02	0.29	8.96	0.07	1.09	4.02	0.29	8.96	0.02	1.57	2.51	0.00	6.98
2.5	0.08	1.13	3.75	0.10	9.21	0.08	1.14	3.78	0.12	9.20	0.06	1.18	3.91	0.22	9.09	0.04	1.44	3.06	0.00	8.30
3.0	0.10	0.94	3.72	0.04	9.24	0.07	1.09	3.58	0.03	9.15	0.09	0.86	3.90	0.16	9.12	0.06	1.18	3.31	0.00	8.75
3.5	0.10	0.87	3.72	0.01	9.30	0.10	0.83	3.80	0.06	9.25	0.09	0.86	3.82	0.10	9.19	0.08	0.97	3.48	0.00	9.01
4.0	0.12	0.79	3.77	0.00	9.33	0.11	0.80	3.79	0.02	9.33	0.10	0.81	3.80	0.06	9.25	0.10	0.88	3.65	0.00	9.24
4.5	0.12	0.74	3.79	0.00	9.34	0.12	0.76	3.77	0.00	9.33	0.11	0.77	3.78	0.04	9.26	0.11	0.78	3.73	0.00	9.29
5.0	0.13	0.72	3.82	0.00	9.35	0.13	0.72	3.81	0.00	9.35	0.12	0.74	3.80	0.02	9.32	0.13	0.73	3.81	0.00	9.35

Table 10

Cost and revenue components (in units of 1000 euros) for all policies when $g = 20$.

ϕ	Optimal					(S, s, v, k)					(S, s, v, k = 0)					(S, s, v, k = M)				
	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R	H	O	P	E	R
1.0	0.01	2.00	3.30	0.26	7.85	0.01	2.00	3.30	0.26	7.85	0.01	2.00	3.30	0.26	7.85	0.00	0.00	0.00	0.00	0.00
1.5	0.03	2.00	3.37	0.10	8.73	0.03	2.00	3.47	0.15	8.70	0.01	2.00	3.30	0.26	7.85	0.01	2.00	1.65	0.00	4.63
2.0	0.05	1.42	3.41	0.04	8.88	0.05	1.09	3.42	0.14	8.37	0.05	1.09	3.42	0.14	8.37	0.02	1.57	2.51	0.00	6.98
2.5	0.07	1.27	3.51	0.01	9.14	0.06	1.19	3.50	0.06	8.94	0.05	1.28	3.50	0.11	8.80	0.04	1.44	3.06	0.00	8.30
3.0	0.09	1.02	3.61	0.00	9.22	0.07	1.06	3.55	0.03	9.10	0.07	0.92	3.56	0.09	8.82	0.06	1.18	3.31	0.00	8.75
3.5	0.10	0.90	3.68	0.00	9.28	0.08	0.97	3.58	0.01	9.18	0.08	0.92	3.57	0.05	9.00	0.08	0.97	3.48	0.00	9.01
4.0	0.12	0.80	3.76	0.00	9.33	0.11	0.85	3.71	0.00	9.30	0.09	0.89	3.65	0.04	9.16	0.10	0.88	3.65	0.00	9.24
4.5	0.12	0.75	3.79	0.00	9.34	0.12	0.77	3.77	0.00	9.33	0.10	0.82	3.68	0.02	9.20	0.11	0.78	3.73	0.00	9.29
5.0	0.13	0.72	3.81	0.00	9.35	0.13	0.73	3.81	0.00	9.35	0.11	0.78	3.75	0.01	9.29	0.13	0.73	3.81	0.00	9.35

of ϕ increases, the total ordering costs tend to be lower while the holding costs tend to increase.

Regarding the purchasing cost, we observe in Table 8 that the optimal policy tends to have a lower purchasing cost than the (S, s, v, k) and the $(S, s, v, k = 0)$ policies. Yet, the optimal policy generates more revenue because it can make a more effective use of the inventories as this policy induces less removal than the other two policies; the optimal policy can avoid excessive removal because it can better deal with quality issues (i.e., due to its flexibility to choose the amount of LNG to remove and purchase). When comparing the cases where the removal of LNG generates a revenue or a cost in Tables 9 and 10, we observe that the purchasing costs of all policies, excepting the $(S, s, v, k = M)$ policy, are much lower when the removal induces a cost. This occurs because the policies tend to choose a conservative order up to level S in order to avoid quality issues that may lead to the removal of LNG.

When comparing the revenue of the (S, s, v, k) and the $(S, s, v, k = 0)$ policies, we observe that the former has a larger revenue. This is caused by the difference in safety stock between policies; in the $(S, s, v, k = 0)$ quality issues are prevented by considering less safety stock than in the (S, s, v, k) , since quality issues entail the removal of all the on-hand LNG. In regard to the revenue of the $(S, s, v, k = M)$ policy, we see that the revenue is much lower than that of all other three policies, which is a consequence of the conservative approach of this policy to tackle quality issues and, in extreme cases, not sell any LNG.

7.2. Sensitivity analysis of the optimal expected profit

In this section, we aim to gain insights on how the input parameters of the problem affect the optimal profit of the system. *Experimental design.* We created a new set of experiments, using a 2^K factorial design to identify cross-parameter relations. Specifically, we considered two levels per parameter, which are the maximum and the minimum values presented in Table 6. An exception is made for parameter ϕ , for which we tested all values for all instances. Furthermore, we fixed the value of $\rho = 0.1$, $\Delta_I = 0.5$ and $\Delta_W = 0.5$. The results obtained were analyzed graphically. Additional experiments were performed for the cases in which we encountered interesting relations among the input parameters.

Results. The overall effect of each of the input parameters on the profit is intuitive. The profit increases with larger values of μ , ϕ and p , while it decreases with larger values of K , c , h , g and σ . In instances where g is negative (where removed LNG has some value), we observe an increase in profit, especially when the quality criticality is high. This increase is not only driven by the extra revenue obtained when removing LNG, but also because the optimal policy tends to place larger replenishment orders, which in turn decrease the replenishment frequency. The increase in the replenishment order size is a natural consequence of setting g as a revenue because it reduces the penalty of the excess of inventory caused by large order sizes.

Throughout the experiments, we observed an interaction among quality criticality, tank size and the mean of the demand. In instances where the mean of the demand is relatively large with respect to the tank size, the system is less vulnerable to quality criticality (i.e., the profit is relatively stable across different levels of ϕ). This is due to the fact that a relatively small tank size induces frequent replenishment, which means that the quality of the LNG in stock is often upgraded based on an inventory level threshold as opposed to a quality threshold. As a result, the profit is less sensitive to changes in quality criticality. To better illustrate this, Fig. 6a shows the profit of the system across all values of ϕ for three different levels of α , which we define as the percentage ratio between the mean of the demand and the tank size. In the cases in which $\phi = 1$ and $\phi = 1.5$, the value of α does not make a dif-

ference in the performance because in these cases a replenishment order is needed in every period. However, for the other values of ϕ , α clearly influences the extent to which the profit is affected by quality criticality.

Another relevant interaction between input parameters occurs between quality criticality and the standard deviation of the demand. Fig. 6b shows the profit across different values of ϕ for three different levels of the standard deviation of the demand σ . The figure indicates that a large standard deviation is more detrimental for the profit when the quality criticality is high. This pattern is caused by the increased amount of LNG that needs to be removed from the system when an inventory control policy increases its safety stock, which is a natural response to a large standard deviation of the demand. Accordingly, as the quality criticality decreases, the amount of LNG that is expected to be removed from the system tends to be lower, which mitigates the effect of a high standard deviation of the demand. Note that there are other parameters that affect the relation between ϕ and σ such as the removal cost g and the unit cost c . Specifically, when the sum of $g + c$ is a small value, the economic loss of removing a unit of LNG is low, which mitigates the effect of σ in the profit when the criticality is high.



8. Conclusions and future research

In this paper, we studied the LNG inventory control problem that arises in LNG storage and refueling facilities. The challenge of managing LNG inventories is that they are subject to continuous boil-off, which results in quantity decay and quality deterioration of the LNG in stock. Since LNG end-users enforce a minimum quality requirement, quality is an essential element of the problem. A unique aspect of the problem is that LNG is a mixable liquid. As a result, the quality of the LNG in stock can be upgraded by means of replenishment with “fresh” LNG. The removal of off-spec LNG from the system is no longer a necessity, which entails that it is a decision variable of an inventory policy for this problem.

We used a Markov decision process model to obtain optimal policies for different scenarios. By means of illustrative examples, we gained insights into the behavior of optimal policies. These insights were later used to design the (S, s, v, k) policy. In a numerical study, we showed that the difference in performance of the (S, s, v, k) policy with respect to the optimal was 1.65% on average. However, in instances where quality criticality is low, the (S, s, v, k) performed close to optimality, since in those cases the optimal policy is purely driven by inventory levels where an (S, s) policy is optimal.

The main conclusions of this study are threefold. Firstly, our study shows that it is important to take quality considerations into account when designing inventory policies for LNG. When inventory management responses are purely driven by quantity triggers, such as in (S, s) policies, the inventory system might fall into a temporal deadlock situation, i.e., where the on-hand LNG can neither be used to serve demand (due to the low quality of the LNG) nor upgraded to meet the minimum quality requirement (due to the lack of response of the policy to the quality of the LNG).

Secondly, we showed that inventory policies that use both removal and mixing of LNG as mechanisms to cope with quality issues, perform better than those policies that exclusively use removal or mixing. However, when only one of the mechanisms is considered, a policy that does not allow removal on average performs worse than a policy where all LNG is removed when quality falls below a given threshold. It appears that the ability to remove on-hand LNG is important for the inventory policy because (1) it prevents the system from falling into a temporal dead-lock situation and (2) it can reduce the number of replenishment orders that need to be performed as result of low-quality LNG. An important

(a) Profit VS ϕ for different levels of α (b) Profit VS ϕ for different levels of σ Fig. 6. Sensitivity of the profit with respect to σ and α for various levels of ϕ .

managerial implication is that LNG storage and refueling facilities can benefit from implementing alternative uses for the removed LNG. For example, the removed LNG can be vaporized and used to produce electricity, or compressed into Compressed Natural Gas and sold as a fuel.

Thirdly, we showed that quality criticality (i.e., the remaining time that the on-hand LNG remains above the minimum quality requirement) has a considerable effect on the profit of the LNG facility. In supply chains where LNG quality criticality is inherently high, for example when the source of LNG is of limited quality, it may be worthwhile to invest in better-insulated storage tanks to decelerate the quality deterioration. Alternatively, one may consider sourcing LNG from another supplier with a higher quality LNG. Producers of bio-LNG can become key players in the small-scale LNG supply chain, as bio-LNG has a very high-quality level, which could be used to considerably upgrade the quality of the on-hand LNG.

We believe that there are ample opportunities for future research on LNG inventory management. First of all, we assumed in our study that the quality of the LNG from the supplier is deterministic, stationary and known. Hence, it would be interesting to study, for example, the case in which the quality of the LNG from the supplier is stochastic but known before the replenishment decision is made. Another interesting aspect that can be considered is that in LNG markets there might be multiple suppliers, each of which provides LNG with different quality and price. Finally, it could be interesting to study the case in which the LNG facility makes use of dynamic pricing to influence the demand for LNG, which can also be used as a mechanism to prevent quality issues.

Acknowledgments

This research was supported in part by the European Union via the INTERREG VA-project LNG PILOTS (No. 151090). The authors are thankful to Sander Gersen (expert on gas quality and combustion) from DNV GL for his contribution to the paper regarding the physio-chemical properties of LNG.

Appendix A. Discussion of the LNG properties considered in this study

This appendix provides further information about the physio-chemical properties of LNG and a justification for the modeling

choices we made in that regard. Throughout a small-scale LNG supply chain, the fuel is transported and stored in well-insulated storage tanks that minimize the exposure of the LNG to heat leakages. Since there is no perfect insulation, however, the ingress of heat to the storage tank causes continuous vaporization of the fuel [8]. This vaporization not only decreases the quantity of LNG in stock but also changes its quality. Specifically, the composition of the fuel changes (this process is also known as weathering in the LNG industry). Accurately estimating changes in quantity and quality of LNG over time involves complex calculations; hence, we make assumptions in our study regarding the physio-chemical properties of LNG. These assumptions are: (1) the quality deterioration b is constant per unit of time; (2) the quantity decay θ is a constant volume per unit of time; and (3) the quality of a mixture of LNG is the weighted average of the loads involved in the mixture. Below, we elaborate on each of these assumptions.

A1. Quality deterioration

The methane number is a well-established quality measure for LNG as a fuel and is determined by the chemical composition. It serves as an indicator of the fuel's knock resistance. If an engine runs on LNG with a lower knock resistance than given in the engine specification, there is a risk of engine knock (i.e., spontaneous ignition). This not only reduces fuel efficiency but could also damage the engine in extreme cases. As a consequence, manufacturers of LNG engines strongly recommend end-users to ensure that their trucks or ships are fueled with LNG that meets the quality specification of the engine.

Theoretically, the methane number of LNG can range from 0 to 100, where 0 is pure hydrogen and 100 is pure methane. LNG with a methane number of 80, for example, is equivalent to a fuel with a mixture of 80% methane and 20% hydrogen. The methane number thus increases when the concentration of light components of the fuel (e.g., methane and nitrogen) increases with respect to that of heavy hydrocarbons (e.g., ethane and propane). Depending on the region in which the LNG is produced, the composition (and thus the quality) of LNG varies. For example, LNG from Nigeria has a methane number of 71 and that of Norway is 76. In addition to the regional differences, the quality of LNG changes over time when moving through or stored in the small-scale LNG supply chain. This quality change occurs because the lighter components of LNG vaporize preferentially with respect to other heavier

Table A1
composition (%mol) and methane number of different types of LNG.

Type	Nitrogen	Methane	Ethane	Propane	n-Butane	Iso-Butane	Methane Number
Type A	0.10	91.30	4.60	2.60	0.80	0.60	71
Type B	0.00	90.29	5.19	2.93	0.90	0.68	69
Type C	0.00	89.06	5.85	3.31	1.02	0.76	67
Type D	0.00	87.46	6.71	3.79	1.17	0.87	65

components, leading to a relative increase of heavy hydrocarbons of the LNG, and hence a lower methane number.

A2. Quantity decay

The fact that LNG supply chains are confronted with decaying quantity is acknowledged in the Operations Research literature, especially in the context of maritime shipping [13,16]. These studies commonly assume a constant boil-off rate, measured as the relative amount of LNG that evaporates per unit of time with respect to the initial volume. In the context of LNG road transport and inland storage, boil-off has been modeled as a fraction of the on-hand inventory level [10]. In reality, the boil-off rate of LNG is a complex process that strongly depends on the properties of the tank [8]. The estimated amount of boil-off can differ significantly depending on whether the boil-off rate is assumed to be constant or dependent on the inventory level [11].

In the context of above-ground storage tanks, which is the setting of our problem, a model was developed in which the long-term sensitivity of the boil-off rate was tested in different experimental conditions, such as different LNG compositions and inventory levels [17]. The results show that the amount of boil-off generated is constant and independent of the inventory levels at higher LNG quality levels. Only when the LNG composition considerably changes (e.g., due to long-term deterioration), the amount of boil-off generated starts to decrease. Indeed, the LNG compositions common in the small-scale LNG supply chain are well within the constant range of boil-off, i.e., where a constant amount of LNG boils off per unit of time. We thus modeled boil-off accordingly.

A3. LNG mixture

Determining the methane number when mixing two loads of LNG involves solving various complex chemical equations that are dependent on the specific composition and the quantity of the LNG loads involved. Since the scope of our paper is on the inventory control implications of the physio-chemical properties of LNG – and not on those properties themselves – we have made the simplifying assumption that the quality of a mixture of LNG is the weighted average of both loads involved in the mixture. Albeit simplifying, this assumption proved to be reasonably accurate after testing some practical examples using Wärtsilä software [1] to estimate the actual methane number when mixing different loads of LNG as shown below. Table A1 presents the different LNG compositions considered in our example. LNG of Type A closely resembles the composition of LNG produced in Nigeria, which has a quality that is common for the LNG imported to the European market. All the other compositions in Table A1 are deteriorated versions of Type A LNG; that is, with relatively less light components with respect to heavy components. In our examples, the different types of deteriorated LNG are upgraded using the Type A LNG.

Table A2 shows the comparison between the actual methane number (i.e., using Wärtsilä software) and the methane number based on the weighted average (i.e., the assumption made in our paper) when mixing loads of LNG of different compositions and quantities. The results indicate that the weighted average is a good

Table A2
Estimated and actual methane number for different mixtures.

Types of LNG mixed		Volume		Methane Number	
Type 1	Type 2	Type 1	Type 2	Actual	Weighted average
B	A	1	1	70.0	70.0
C	A	1	1	69.0	69.0
D	A	1	1	68.0	68.0
B	A	1	2	70.0	70.3
C	A	1	2	70.0	69.7
D	A	1	2	69.0	69.0
B	A	1	3	70.0	70.5
C	A	1	3	70.0	70.0
D	A	1	3	70.0	69.5

approximation of the actual methane number for loads of LNG considered in this study. Indeed, it is fully accurate for five of the mixtures considered in this study. In the other four mixtures, the error between the actual and the estimated methane number is reasonably small.

References

- [1] Wärtsilä Wärtsilä Methane Number Calculator. <https://www.wartsila.com/products/marine-oil-gas/gas-solutions/methane-number-calculator>. Accessed: 2018-05-30.
- [2] Andersson H, Christiansen M, Desaulniers G. A new decomposition algorithm for a liquefied natural gas inventory routing problem. *Int J Prod Res* 2016;54(2):564–78.
- [3] Axsäter S. *Inventory control*. 225. Springer; 2015.
- [4] Bakker M, Riezebos J, Teunter RH. Review of inventory systems with deterioration since 2001. *Eur J Oper Res* 2012;221(2):275–84.
- [5] Cai X, Chen J, Xiao Y, Xu X. Optimization and coordination of fresh product supply chains with freshness-keeping effort. *Prod Operat Manag* 2010;19(3):261–78.
- [6] Cai X, Chen J, Xiao Y, Xu X, Yu G. Fresh-product supply chain management with logistics outsourcing. *Omega (Westport)* 2013;41(4):752–65.
- [7] Chen J, Dong M, Rong Y, Yang L. Dynamic pricing for deteriorating products with menu cost. *Omega (Westport)* 2018;75:13–26.
- [8] Dobrota Đ, Lalić B, Komar I. Problem of boil-off in LNG supply chain. *Trans Marit Sci* 2013;2(02):91–100.
- [9] Fokkema JE, Buijs P, Vis IF. An investment appraisal method to compare LNG-fueled and conventional vessels. *Trans Res Part D Trans Environ* 2017;56:229–40.
- [10] Ghiami Y, Van Woensel T, Christiansen M, Laporte G. A combined liquefied natural gas routing and deteriorating inventory management problem. In: *Proceedings of the international conference on computational logistics*. Springer; 2015. p. 91–104.
- [11] Głomski P, Michalski R. Problems with determination of evaporation rate and properties of boil-off gas on board LNG carriers. *J Pol CIMAC* 2011;6(1):133–40.
- [12] Goyal S, Giri BC. Recent trends in modeling of deteriorating inventory. *Eur J Oper Res* 2001;134(1):1–16.
- [13] Grønhaug R, Christiansen M, Desaulniers G, Desrosiers J. A branch-and-price method for a liquefied natural gas inventory routing problem. *Transp Sci* 2010;44(3):400–15.
- [14] Inderfurth K. Simple optimal replenishment and disposal policies for a product recovery system with leadtimes. *Operat Res Spektrum* 1997;19(2):111–22.
- [15] Janssen L, Claus T, Sauer J. Literature review of deteriorating inventory models by key topics from 2012 to 2015. *Int J Prod Econ* 2016;182:86–112.
- [16] Jokinen R, Pettersson F, Saxén H. An MILP model for optimization of a small-scale LNG supply chain along a coastline. *Appl Energy* 2015;138:423–31.
- [17] Migliore C, Tubilleja C, Vesovic V. Weathering prediction model for stored liquefied natural gas (LNG). *J Nat Gas Sci Eng* 2015;26:570–80.
- [18] Molana S, Davoudpour H, Minner S. An (R, NQ) inventory model for packaged deteriorating products with compound poisson demand. *J Operat Res Soc* 2012;63(11):1499–507.
- [19] Nahmias S. *Perishable inventory systems*, 160. Springer Science & Business Media; 2011.

- [20] Post RM, Buijs P, uit het Broek MA, Lopez Alvarez JA, Szirbik NB, Vis IF. A solution approach for deriving alternative fuel station infrastructure requirements. *Flex Serv Manuf J* 2018;30(3):592–607.
- [21] Puterman ML. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons; 2014.
- [22] Rajan A, Rakesh, Steinberg R. Dynamic pricing and ordering decisions by a monopolist. *Manage Sci* 1992;38(2):240–62.
- [23] Rakke JG, Stålhane M, Moe CR, Christiansen M, Andersson H, Fagerholt K, et al. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transp Res Part C Emerg Technol* 2011;19(5):896–911.
- [24] Sarkar B, Saren S, Cárdenas-Barrón LE. An inventory model with trade-credit policy and variable deterioration for fixed lifetime products. *Ann Oper Res* 2015;229(1):677–702.
- [25] Thunnissen SK, van de Bunt LG, Vis IF. Sustainable fuels for the transport and maritime sector: a blueprint of the LNG distribution network. In: *Logistics and Supply Chain Innovation*. Springer; 2016. p. 85–103.
- [26] Veinott AF Jr. On the optimality of (s,s) inventory policies: new conditions and a new proof. *SIAM J Appl Math* 1966;14(5):1067–83.
- [27] Wang WC, Teng JT, Lou KR. Seller'S optimal credit period and cycle time in a supply chain for deteriorating items with maximum lifetime. *Eur J Oper Res* 2014;232(2):315–21.
- [28] Wu J, Chang CT, Teng JT, Lai KK. Optimal order quantity and selling price over a product life cycle with deterioration rate linked to expiration date. *Int J Prod Econ* 2017;193:343–51.