Applied Energy 412 (2026) 127670

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apen

Adaptive liquefied natural gas supply chain planning under demand uncertainty

Euihyeon Choi [a,1], Junhyeok Lee [b,1], Ilkyeong Moon [c,d,*]

[a] LG CNS, Seoul 07795, Korea
[b] Department of Industrial and Management Engineering, Incheon National University, Incheon 22012, Korea
[c] Department of Industrial Engineering, Seoul National University, Seoul 08826, Korea
[d] Institute for Industrial Systems Innovation, Seoul National University, Seoul 08826, Korea

HIGHLIGHTS

- Nationwide LNG supply chain optimization under demand uncertainty.
- A two-phase approach ensures near-optimal performance under uncertainty.
- The model enables proactive planning and reactive responses to realized demand.
- Demonstrates a 5% optimality gap against the perfect information lower bound.
- Sensitivity analysis reveals key insights for LNG supply chain management.

ARTICLE INFO

Keywords:
Liquefied natural gas
Demand uncertainty
Supply chain planning
Adaptive decision-making
Robust optimization

ABSTRACT

Global liquefied natural gas trade volumes have risen sharply as countries seek cleaner energy alternatives. As liquefied natural gas emerges as a primary energy source in many regions, securing it efficiently has become strategically critical. However, supply plans based on point forecasts often become infeasible or excessively costly when actual demand deviates from forecasts. To address this challenge, we propose a new optimization framework that explicitly accounts for demand uncertainty and adaptively updates decisions as real-world demand unfolds. We formulate the nationwide liquefied natural gas supply chain planning problem as a stochastic optimization model that jointly optimizes long-term contract selection, spot market procurement, volume allocation, and inventory control. To overcome the computational intractability of this formulation, we develop an adaptive robust optimization framework solved via a customized two-phase approach. Large-scale computational experiments calibrated with data from South Korea's liquefied natural gas supply chain demonstrate that the proposed framework achieves an optimality gap within 5% relative to a perfect information lower bound, substantially outperforming the industry-standard deterministic benchmark. Furthermore, through extensive sensitivity analysis, we derive critical managerial insights for energy planners, identifying spot price as the dominant cost driver. The proposed framework offers national energy planners a practical and robust decision-support tool for liquefied natural gas supply chain planning under demand uncertainty.

1. Introduction

Liquefied Natural Gas (LNG) is natural gas that is cooled to approximately -162 °C, transforming it from a gaseous state into a liquid state. This liquefaction reduces its volume to about 1/600th of its original state, greatly facilitating storage and transportation over long distances. This is especially important for countries that cannot import natural gas via pipeline networks due to geopolitical issues. As a result, LNG has become a key energy source. It has also played a significant role in changing global natural gas trade patterns.

Over the past two decades, global trade in LNG has grown substantially. In 2024, global LNG trade reached approximately 407 million tonnes with a compound annual growth rate of 6.3% over the past two

* Corresponding author at: Department of Industrial Engineering, Seoul National University, Seoul 08826, Korea.
Email addresses: ceh9856@snu.ac.kr (E. Choi), ljh9533@inu.ac.kr (J. Lee), ikmoon@snu.ac.kr (I. Moon).

[1] These authors contributed equally to this work.

https://doi.org/10.1016/j.apenergy.2026.127670
Received 4 December 2025; Received in revised form 4 March 2026; Accepted 7 March 2026
Available online 16 March 2026
0306-2619/

decades [1,2]. This growth is primarily due to the environmental efforts of countries worldwide. LNG is becoming more attractive because it emits fewer carbon emissions than coal when combusted. Countries that aim to reduce greenhouse gas emissions are increasingly using LNG as a transitional energy source.

Mainly, three countries in East Asia—China, Japan, and South Korea—account for nearly half of the global LNG imports, underscoring the strategic importance of LNG supply chain planning for energy security in the region. Specifically, South Korea is entirely dependent on imports to meet its natural gas demand, with annual import volumes ranging from 45 to 50 million tonnes. Consequently, efficient and cost-effective management of LNG procurement has significant national economic implications for these countries.

Because of the significant capital investments involved in establishing LNG production and export facilities, international LNG trade has historically been dominated by *long-term contracts* (LTCs) [3]. These contracts typically span multiple years, often exceeding five years, and are secured in advance to guarantee large-volume supplies at predetermined prices. LTCs provide supply stability and typically offer lower unit costs compared to alternative procurement options; however, they offer limited flexibility in volume adjustments. As LNG demand inherently fluctuates due to seasonal and economic factors, solely relying on fixed-volume LTCs becomes impractical for fully addressing consumer demand. This limitation necessitates supplementary procurement from the LNG spot market, where cargoes are traded for immediate delivery without LTCs.

The *spot market* has expanded significantly in recent years, driven by increased LNG production capacity, market liberalization, and a growing preference among buyers for contractual flexibility [4]. By allowing buyers to adjust their procurement quantities and timing more responsively, spot market procurement helps mitigate the challenges posed by seasonal demand peaks and unforeseen short-term fluctuations. Nevertheless, LNG procured through the spot market generally involves higher per-unit costs and greater exposure to price volatility [5]. Thus, LNG procurement strategies must carefully balance the trade-off between the cost efficiency of LTCs and the operational flexibility offered by the spot market.

Effective LNG sourcing and procurement planning thus hold substantial national economic significance. Because of the high infrastructure and operational costs associated with LNG—including specialized cryogenic storage terminals, dedicated LNG vessels, and substantial inventory holding expenses—the effectiveness of supply planning decisions directly impacts overall procurement expenditures [6]. Countries heavily reliant on LNG imports, such as Japan and South Korea, often establish long-term natural gas procurement strategies that span several decades. These strategies have traditionally relied on deterministic demand forecasts to structure LTC portfolios and anticipated spot market procurement. However, inaccuracies in demand forecasting frequently result in either surplus or deficit procurement, substantially increasing procurement costs and operational inefficiencies.

Numerous studies have addressed LNG supply chain planning from various perspectives. Prior work has focused on minimizing procurement costs through optimal contract selection [7], managing supply risks and security [8], and integrating spot market flexibility with long-term decisions [9]. Despite these advances, existing approaches face critical limitations. One major limitation is that scenario-based models can be challenging to implement when actual demand deviates from predefined scenarios, limiting their practical applicability. Another critical issue lies with forecast-driven models, which often rely on the unrealistic assumption that decisions for future periods are made without incorporating the realized demand information available at the decision time. For example, a spot market procurement decision made five years in advance cannot adaptively utilize demand observations from the preceding years.

To address these limitations, we develop an adaptive robust optimization (ARO) framework for nationwide LNG supply chain planning. In contrast to existing scenario-based and forecast-driven approaches, our framework makes decisions without restrictive assumptions about demand uncertainty and adaptively incorporates realized demand information. The proposed framework makes five key contributions. First, we formulate the problem as a multistage stochastic model that jointly optimizes LTC selection, spot procurement, allocation, and inventory control. Second, we develop a tractable two-phase solution approach that combines target-oriented robust optimization (TRO) with linear decision rules (LDRs), enabling policies that adapt dynamically as demand is realized. Third, the framework proactively accounts for demand forecast errors while exploiting realized demand information at each decision point, supporting resilient and economical procurement plans for national energy planners. Fourth, we demonstrate practical applicability through large-scale experiments calibrated with South Korean data, achieving near-optimal performance with a 5% gap from perfect information. Fifth, beyond methodological contributions, this study offers practical managerial insights derived from rigorous sensitivity analyses. We identify key cost drivers and provide strategic guidelines on balancing LTCs with spot market flexibility, thereby serving as practical guidelines for policy makers.

The paper proceeds as follows. Section 2 reviews the literature on LNG supply chain planning and robust optimization (RO), and outlines the contributions of this study. Section 3 formally defines the nationwide LNG supply chain problem and presents a stochastic optimization model that incorporates demand uncertainty. Because the resulting model is computationally intractable in its original formulation, Section 4 introduces a customized two-phase solution procedure that combines a TRO in Phase 1 with LDR in Phase 2, which is called TPA. Section 5 reports comprehensive numerical experiments that evaluate the proposed approach. These results offer technical insights and practical guidelines for governments and procurement planners in countries that heavily rely on imported LNG. Finally, Section 6 summarizes the main findings and suggests directions for future research.

2. Literature review

2.1. LNG supply chain planning and procurement

Previous studies have developed optimization-based approaches for LNG supply chain planning, addressing contract selection, risk management, and operational flexibility. Khalilpour and Karimi [7] pioneered a mathematical programming approach aimed at minimizing LNG procurement costs by optimally selecting among available LTCs. Biresselioglu et al. [10] proposed structured optimization models to manage supply risks effectively, highlighting their implications for national energy security. Additionally, Geng et al. [8] emphasized the need to incorporate multiple risk factors—including market price volatility, geopolitical considerations, and supply security—when designing optimal LNG procurement portfolios.

More recently, Sangaiah et al. [11] applied RO methodologies to LNG supply chain sales planning, addressing manufacturer supply uncertainty. Shahrukh et al. [9] further advanced LNG procurement methodologies by integrating spot market operational flexibility with LTC selection decisions using mathematical programming. The authors emphasized the advantages of incorporating dynamic decisions in supply chain planning. Eriksen et al. [12] investigated multistage stochastic programming approaches for maritime LNG supply chain planning under demand uncertainty, demonstrating the value of capturing uncertainty evolution over time through scenario trees. Al-Haidous et al. [13] integrated sustainability and resilience considerations into LNG supply chain optimization, using a Qatari case study. This study illustrated the trade-offs between economic efficiency, environmental impact, and supply chain robustness. Li et al. [14] addressed the maritime inventory routing problem with transshipment operations in the context of the Yamal LNG project. Furthermore, the authors proposed solution methods that account for the complex logistical coordination inherent in LNG maritime supply chains.

Despite these contributions, existing approaches predominantly rely on deterministic demand forecasts or predefined scenario sets, limiting their ability to incorporate realized demand information adaptively over the planning horizon. This limitation motivates the development of ARO frameworks that can dynamically adjust decisions as uncertainty is resolved.

2.2. Robust optimization

RO provides a framework for decision-making under uncertainty without requiring complete knowledge of the probability distributions governing uncertain parameters [15]. Unlike stochastic programming, which relies on distributional assumptions, RO assumes that uncertain parameters belong to a predetermined *uncertainty set*—typically defined as a convex region such as polyhedral [16] or ellipsoidal sets [17]. The objective of RO is to identify solutions that remain feasible and perform well under the worst-case realization within this uncertainty set.

For multistage decision problems, two fundamental decision paradigms exist: *here-and-now* decisions, which are determined before uncertainty is revealed, and *wait-and-see* decisions, which can be adjusted after observing partial or complete information about uncertain parameters. The latter approach, known as ARO, offers less conservative solutions by allowing decisions to adapt as uncertainty unfolds [18]. However, ARO models are often computationally intractable due to the large feasible space of adjustable variables.

To address tractability challenges, ARO commonly employs *decision rules*—parameterized functions that prescribe how adjustable decisions respond to realized uncertainties. The LDRs, introduced by Ben-Tal et al. [19], express adjustable variables as affine functions of uncertain parameters and have been widely adopted in inventory management and production planning [20,21]. While nonlinear decision rules have been explored [22,23], LDRs remain popular due to their computational tractability and strong empirical performance. In certain problem structures, even simpler *static rules*—where decisions remain fixed regardless of realized uncertainty—have been proven optimal [24–26].

Extending ARO to problems involving binary adjustable variables presents additional challenges. Several methodologies have been developed to address this issue, including K-adaptability [27], finite adaptability [28,29], and binary decision rules [30]. Among these approaches, Lim and Wang [25] introduced TRO, a goal-oriented framework that seeks to maximize the size of the uncertainty set for which a prespecified cost target can be achieved while simultaneously handling both binary and continuous adjustable variables.

These decision rule methodologies, from static rules to LDRs, have significantly advanced the tractability of continuous ARO problems. However, extending ARO to binary adjustable variables remains a fundamental challenge. TRO offers a promising framework for jointly optimizing binary and continuous adjustable variables through its goal-oriented structure, yet its application to large-scale, multi-echelon supply chain problems with complex operational constraints has not been explored.

2.3. Adaptive robust optimization in supply chain planning and energy procurement

ARO has been successfully applied to various supply chain and procurement planning problems, demonstrating its effectiveness in managing demand uncertainty while maintaining operational feasibility. Ben-Tal et al. [31] developed an RO framework for retailer–supplier flexible commitment contracts, showing that affinely adaptive policies can achieve near-optimal cost performance under multi-period demand uncertainty. Kim and Chung [32] applied ARO to multi-period, multi-product production planning, demonstrating that LDRs provide significantly less conservative solutions compared to static policies while maintaining worst-case feasibility guarantees.

Building upon the TRO framework, Lim et al. [33] proposed the TPA for online retail operations. The TPA decouples the optimization of binary and continuous adjustable variables: binary decisions (such as facility activation) are determined through static rules derived from TRO, while continuous decisions (such as allocation and fulfillment quantities) are optimized using LDRs. This decomposition strategy enables the model to handle large-scale problems efficiently while achieving near-optimal performance. Their computational results demonstrated that TPA outperforms alternative methodologies in both solution quality and computational efficiency. More recently, researchers have extended RO techniques to omnichannel retail operations, effectively balancing inventory positioning and fulfillment decisions under complex, multi-channel demand uncertainty [34,35].

Regarding energy procurement, ARO has emerged as a prominent methodology for addressing supply uncertainty and infrastructure planning challenges. Riepin et al. [36] deployed an ARO framework for European gas infrastructure expansion planning, addressing long-term uncertainties on both the supply and demand sides. For integrated energy systems, Gu et al. [37] applied RO to power-to-gas management, demonstrating how hydrogen converted from wind power can be injected into natural gas systems while maintaining gas quality under supply uncertainty. Shen et al. [38] introduced a data-driven RO framework for large-scale industrial energy systems, where the uncertainty set is constructed using support vector clustering. In power system planning, Abdin et al. [39] developed a multistage ARO model for generation expansion planning. The proposed plans account for short-term operational constraints while guaranteeing the non-anticipativity of uncertainty realizations. Bazyar et al. [40] proposed an RO framework for natural gas supply chain networks under uncertainty, incorporating possibilistic programming methods to handle parameter uncertainty in infrastructure design. For renewable energy systems, adaptive distributionally RO approaches have been applied to hybrid renewable energy system sizing, effectively hedging against uncertainty in renewable generation and demand patterns [41].

While these applications demonstrate the broad applicability of ARO across supply chain and energy domains, no existing study addresses the joint optimization of all four key LNG procurement decisions—LTC selection, spot market procurement, terminal allocation, and inventory management—within a unified adaptive framework. This research gap is particularly significant for national-scale LNG systems, where these decisions are tightly coupled and must be coordinated under demand uncertainty over long planning horizons.

2.4. Research gaps and contributions

Despite these methodological advances, significant research gaps remain in applying ARO to nationwide LNG supply chain planning. While previous studies focused on mathematical modeling, few have explicitly analyzed how variations in key market parameters (e.g., spot prices, supply amounts) alter optimal procurement strategies. The present study addresses these gaps and makes the following contributions:

- We formulate an ARO model for nationwide LNG supply chain planning that jointly optimizes LTC selection, spot market procurement, terminal allocation, and inventory management. The proposed model captures the full complexity of strategic LNG procurement decisions under demand uncertainty.
- Unlike existing LNG planning approaches that rely on deterministic forecasts or predefined scenario trees, our framework adaptively incorporates realized demand information at each decision point, enabling policies that adjust dynamically as uncertainty unfolds over the planning horizon.
- We develop a tractable two-phase solution approach, called TPA, that integrates TRO with LDRs. The proposed approach allows

Table 1
Comparison of related studies on LNG supply chain planning and robust optimization.

Author(s)	LTC selection	Spot market	Allocation	Inventory	Uncertainty	Solution method
Ben-Tal et al. [31]	✓	✓		✓	Robust	Affine
Bertsimas & Thiele [20]				✓	Robust	Robust
Khalilpour & Karimi [7]	✓		✓	✓		MILP
Geng et al. [8]	✓	✓			Stochastic	NSGA-II[a]
Lim & Wang [25]				✓	Robust	TRO
Sangaiah et al. [11]	✓			✓	Robust	Robust MILP
Shahrukh et al. [9]	✓	✓	✓	✓		MILP
Eriksen et al. [12]		✓	✓	✓	Stochastic	MSP[b]
Al-Haidous et al. [13]			✓	✓		MILP + BPSO[c]
Riepin et al. [36]		✓	✓		Robust	CCG[d]
Abdin et al. [39]			✓		Robust	Multistage[e]
This study	✓	✓	✓	✓	**Robust**	**TRO + LDR**

[a]Non-dominated sorting genetic algorithm II; [b]Multistage stochastic programming; [c]MILP with binary particle swarm optimization; [d]Column-and-constraint generation; [e]Multistage ARO.

for the simultaneous optimization of both binary contract selection decisions and continuous operational decisions while maintaining computational tractability for large-scale, long-horizon problems.

- Through comprehensive computational experiments calibrated with South Korean data, we demonstrate that the proposed approach achieves near-optimal performance within a 5% optimality gap relative to perfect information bounds, while outperforming the various benchmark policies.
- Finally, we perform an extensive sensitivity analysis to distill practical managerial insights. This allows us to offer concrete policy recommendations on how to adjust procurement strategies in response to external market changes.

Table 1 summarizes the key characteristics of related studies, including the solution method employed (e.g., mixed-integer linear programming (MILP)). While previous research has addressed various aspects of LNG supply chain planning and RO, our work is the first to jointly consider all four key decisions—LTC selection, spot market procurement, terminal allocation, and inventory management—within an ARO framework.

3. Mathematical models

In this section, we define a nationwide LNG supply chain planning problem and derive the mathematical model required to optimize it. We begin by establishing the notation used throughout our formulations. Tables A.6 and A.7 in Appendix A provide a complete summary of all indices, sets, parameters, and decision variables. The system includes LTC candidates, indexed by $i = 1, \dots, I$, which must be signed before the start of a multi-period planning horizon to secure a stable supply over the years. In addition, the system has a set of specialized LNG terminals, indexed by $j = 1, \dots, J$, each capable of receiving and storing LNG in its liquid state. Demands are realized at each terminal in every period, and the decision maker aims to fulfill this demand using (i) supplies from LTCs and (ii) on-hand inventory. In the event of a supply shortage, the decision maker either purchases additional LNG on the spot market, typically at a premium, or incurs a significant penalty cost due to stockouts.

The time horizon is discretized into periods $t = 1, \dots, T$. While individual periods are discretized into short intervals (monthly or bimonthly), the planning horizon spans five years to fully capture the dynamics of LTCs. The decision maker proceeds in a stage-by-stage fashion over these T periods as follows:

1. *Contract selection (period 0).* Choose the LTC portfolio before operations start. Each LTC specifies different supply terms.
2. *Allocation.* At the start of each period t, the decision maker allocates the contract supply across terminals based on forecasted (nominal) demand, storage capacity, and cost considerations.
3. *Demand realization and inventory usage.* Demand arises at each terminal j. The decision maker meets this demand using on-hand inventory (carried over from previous periods) plus the newly allocated supply in period t.
4. *Spot market procurement and stockout.* If the combination of on-hand inventory and LTC supply is insufficient to meet terminal demand, additional LNG can be procured from the spot market, or a stockout (lost sales) may occur depending on the cost-effectiveness.
5. *Inventory update.* Any excess LNG is retained in storage (subject to each terminal's storage capacity and operational requirements) for use in subsequent periods.

To formulate our model, we define the following sets for indexing decisions and constraints:

$$\mathcal{I} = \{1, \dots, I\}, \quad \mathcal{J} = \{1, \dots, J\}, \quad \mathcal{T} = \{1, \dots, T\}, \quad \mathcal{T}^+ = \{1, \dots, T+1\}.$$

Here, $\mathcal{T}^+$ is used to track end-of-period inventory states through $t = T+1$.

3.1. Deterministic optimization model

We begin by examining a deterministic optimization model in which all demands are known with certainty for each period. The decision-maker aims to minimize total cost by determining (1) which LTCs to sign, (2) how to allocate LTC supply, (3) how to procure from the spot market, and (4) how to manage inventory across periods.

Fig. 1 illustrates the overall structure of the LNG supply chain planning problem. Let x_i for $i \in \mathcal{I}$ be a binary decision variable indicating whether the i-th LTC candidate is signed ($x_i = 1$) or not ($x_i = 0$) prior to the start of the planning horizon. Among the various candidate contracts, each specifies different supply terms:

- *Periodic supply amount,* denoted SUP_i.
- *Contract start and end periods,* START_i and END_i. During periods t within $\mathrm{START}_i \le t \le \mathrm{END}_i$, the supply from contract i is $U_{it} = \mathrm{SUP}_i$, and otherwise $U_{it} = 0$.
- *Unit price,* PRICE_i, the agreed cost per tonne for every shipment delivered under contract i.
- *Fixed cost,* FC_i, covering obligations, shipping premiums, and other contract-specific charges.

Therefore, each contract i entails a total cost as follows:

$$C_i = \mathrm{FC}_i + \mathrm{SUP}_i \times \mathrm{PRICE}_i \times \left(\mathrm{END}_i - \mathrm{START}_i + 1\right),$$

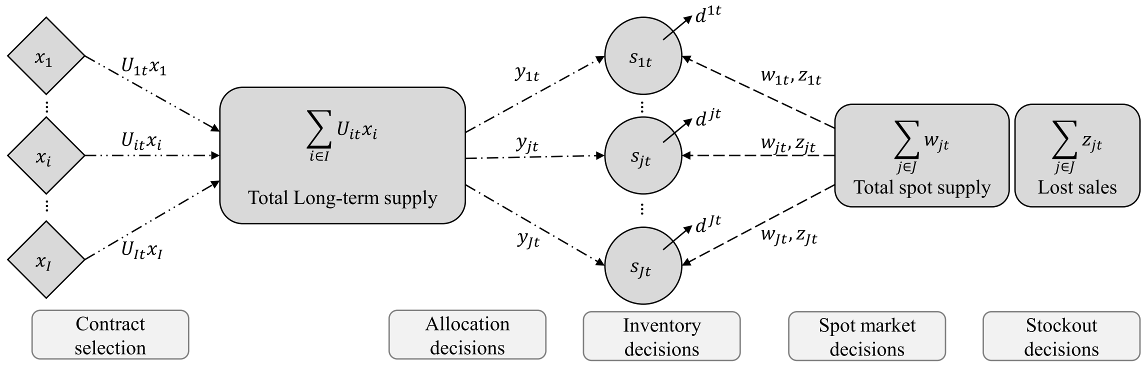

Fig. 1. LNG supply chain network diagram for proposed problem.

which represents the cost of purchasing SUP_i tonnes each period over the contract duration. Using x_i as the signing decision, the total LTC cost over the planning horizon can be expressed as $\sum_{i\in\mathcal{I}} C_i\, x_i$.

Once the planning horizon begins, the total long-term supply available in each period t is $\sum_{i\in\mathcal{I}} U_{it}\, x_i$. This quantity represents the aggregate volume of LNG delivered through all active LTCs. Because each terminal $j \in \mathcal{J}$ has its own demand and operational requirements, the decision-maker must allocate this supply appropriately. Let y_{jt} denote the amount of LNG delivered to terminal j in period t. To ensure that all available supply in period t is fully distributed, we impose: $\sum_{i\in\mathcal{I}} U_{it}\, x_i \;=\; \sum_{j\in\mathcal{J}} y_{jt}, \quad \forall t \in \mathcal{T}$. In other words, no portion of the contracted supply is left unassigned at the end of each period. Transporting LNG to a terminal may incur additional handling costs or short-haul transport costs. Let a_{jt} denote the per-tonne allocation cost for delivering LNG to terminal j in period t. Consequently, the total allocation cost in period t is $\sum_{j\in\mathcal{J}} a_{jt}\, y_{jt}$, and summing over all periods yields the overall allocation expense across the entire planning horizon. This allocation mechanism ensures that every terminal's portion of long-term supply is carefully distributed to balance operational needs, terminal capacities, and cost considerations.

Let s_{jt} denote the initial on-hand inventory at terminal j at the start of period t for $j \in \mathcal{J}, t \in \mathcal{T}^+$. After the decision maker allocates supply to terminal j in period t, the terminal's inventory is temporarily updated to $s_{jt} + y_{jt}$. Because each terminal j can only store LNG up to its maximum capacity $\bar{s}_j$, we impose the constraint $s_{jt} + y_{jt} \;\le\; \bar{s}_j, \quad \forall j \in \mathcal{J},\ \forall t \in \mathcal{T}$. Shortly after this inventory update, demand d^{jt} occurs at terminal j in period t. To meet d^{jt}, the decision maker uses a portion of the on-hand inventory, denoted r_{jt}. If demand exceeds available inventory, the decision maker procures additional LNG from the spot market. However, if the required quantity exceeds the spot market capacity, the unmet portion is treated as lost sales.

We define δ_{jt} as a binary decision variable that indicates whether the decision maker chooses to procure LNG from the spot market at terminal j in period t. Conditional on this decision, let w_{jt} denote the corresponding decision variable representing the quantity of LNG procured from the spot market at terminal j in period t. To ensure that spot procurement occurs only if the spot market is activated, we impose the following constraint: $w_{jt} \;\le\; M\,\delta_{jt}, \quad \forall j \in \mathcal{J},\ \forall t \in \mathcal{T}$, where M represents the maximum possible quantity of LNG procured from the spot market. Whenever the decision maker decides to procure LNG from the spot market in period t for terminal j, a fixed cost f_{jt} is incurred. This cost accounts for expenses such as chartering, contracting, and insurance. If $\delta_{jt} = 1$, a fixed cost f_{jt} is incurred, along with a variable cost $p_t w_{jt}$. Accordingly, the total spot market cost incurred in period t is given by $\sum_{j\in\mathcal{J}} \left(f_{jt}\,\delta_{jt} + p_t\, w_{jt}\right)$. Note that this modeling choice assumes spot procurement lead times can be absorbed within each planning period. In our bi-monthly framework, each period spans approximately 60 days, which is sufficiently long to accommodate typical LNG voyage times to Northeast Asia (17–34 days depending on the origin [42]). This ensures that LNG procured from the spot market during a given period can be delivered and utilized within the same period.

The realized demand at terminal j in period t must be balanced through inventory usage, spot market procurement, and potential stockouts. Let z_{jt} denote the stockout quantity at terminal j in period t, which occurs when the available supply cannot fully meet demand. We adopt a lost sales model for LNG supply chain planning because natural gas demand is inherently time-sensitive: end consumers require immediate energy for heating, power generation, and industrial processes that cannot be deferred. When a terminal fails to meet demand, customers typically switch to alternative energy sources or curtail operations rather than wait for delayed LNG deliveries.

The demand balance equation is defined as follows: $r_{jt} + w_{jt} + z_{jt} \;=\; d^{jt}, \quad \forall j \in \mathcal{J},\ \forall t \in \mathcal{T}$. This formulation allows for unmet demand when supply constraints are binding, with $z_{jt} > 0$ indicating a stockout situation. Stockouts incur a penalty cost l_{jt} per tonne, capturing revenue loss, contractual penalties, and reputational damage from supply disruptions. Finally, the total stockout cost in period t is expressed as $\sum_{j\in\mathcal{J}} l_{jt}\, z_{jt}$. At the end of period t, the remaining inventory at terminal j is given by $(1-\beta_j)s_{jt} + y_{jt} - r_{jt}$, where the factor $(1-\beta_j)$ accounts for boil-off gas (BOG) losses. The $\beta_j \in [0,1)$ denotes the periodic BOG loss rate at terminal j, representing the fraction of stored LNG that evaporates per period. This value becomes the initial inventory $s_{j,t+1}$ for the subsequent period $t+1$. In addition, each terminal j must maintain a minimum inventory level $\underline{s}_j$ to ensure safe and reliable operations under cryogenic temperatures. Thus, we impose $\underline{s}_j \;\le\; s_{jt}, \quad \forall j \in \mathcal{J},\ \forall t \in \mathcal{T}$.

Storing LNG incurs notably high costs [43], primarily because it must be kept at extremely low temperatures to remain in liquid form. Maintaining such cryogenic conditions at terminals requires specialized equipment and continuous cooling, which significantly increases the cost of LNG inventory compared to many other commodities. Furthermore, LNG is susceptible to natural evaporation (BOG). The reliquefaction process is both technically complex and costly. Therefore, it is crucial to account for inventory holding costs in this model explicitly. Let h_t denote the per-tonne holding cost of storing LNG in period t, and define s_{jt+1} as the carried over inventory at terminal j from period t to $t+1$. Then, the total holding cost incurred in period t is given by $\sum_{j\in\mathcal{J}} h_t\, s_{jt+1}, \quad \forall t \in \mathcal{T}$. Because inventory costs have a considerable impact on the total cost, the decision- maker must consider the trade-off between maintaining sufficient inventory to enhance supply reliability and keeping inventory levels low to improve efficiency.

The objective is to minimize the total cost over the planning horizon by determining a set of decisions, including contract selection $\boldsymbol{x} \;=\; \left(x_i, \forall i \in \mathcal{I}\right)$, supply allocation $\boldsymbol{y} \;=\; \left(y_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$, stockout quantities $\boldsymbol{z} \;=\; \left(z_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$, inventory flows $\boldsymbol{r} \;=\; \left(r_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$ and $\boldsymbol{s} \;=\; \left(s_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}^+\right)$, as well as spot market usage $\boldsymbol{\delta} \;=\; \left(\delta_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$ and procurement quantities $\boldsymbol{w} \;=$

$\left(w_{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$. The following deterministic optimization model, $\mathrm{P_{DET}}$, formalizes the relationships above and incorporates all relevant components of the problem. All cost parameters enter the objective function as per-unit rates that scale linearly with quantities. Specifically, these include allocation costs a_{jt}, spot prices p_t, holding costs h_t, and stockout penalties l_{jt}. This linear cost structure is standard in strategic LNG supply chain planning [7,9,11,31], where each parameter represents a contractually fixed or period-averaged unit rate at the bi-monthly planning granularity:

$(\mathrm{P_{DET}})$

$$\min \quad \underbrace{\sum_{i\in\mathcal{I}} C_i x_i}_{\text{LTCs}} + \sum_{t\in\mathcal{T}} \left(\underbrace{\sum_{j\in\mathcal{J}} a_{jt} y_{jt}}_{\text{Allocation}} + \underbrace{\sum_{j\in\mathcal{J}} f_{jt}\delta_{jt}}_{\text{Spot fixed}} + \underbrace{\sum_{j\in\mathcal{J}} p_t w_{jt}}_{\text{Spot procurement}} + \underbrace{\sum_{j\in\mathcal{J}} l_{jt} z_{jt}}_{\text{Lost sales}} + \underbrace{\sum_{j\in\mathcal{J}} h_t s_{jt+1}}_{\text{Holding}} \right) \tag{1}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{I}} U_{it} x_i = \sum_{j\in\mathcal{J}} y_{jt}, \qquad \forall t \in \mathcal{T} \tag{2}$$

$$s_{jt} + y_{jt} \le \bar{s}_j, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{3}$$

$$\underline{s}_j \le s_{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}^+ \tag{4}$$

$$s_{jt+1} = (1-\beta_j)s_{jt} + y_{jt} - r_{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{5}$$

$$r_{jt} + w_{jt} + z_{jt} = d^{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{6}$$

$$w_{jt} \le M\delta_{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{7}$$

$$y_{jt} \ge 0, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{8}$$

$$w_{jt} \ge 0, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{9}$$

$$r_{jt} \ge 0, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{10}$$

$$z_{jt} \ge 0, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{11}$$

$$x_i \in \{0,1\}, \qquad \forall i \in \mathcal{I}, \tag{12}$$

$$\delta_{jt} \in \{0,1\}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{13}$$

We note that $\mathrm{P_{DET}}$ focuses on the core procurement and inventory decisions and does not explicitly include operational constraints such as berth/throughput limits, regasification capacity, and LTC swing clauses. At the strategic planning level with bi-monthly granularity, these constraints are unlikely to be binding for South Korea's LNG system, where terminal receiving capacity (exceeding 100 million tonnes annually [44]) substantially exceeds import volumes (45–50 million tonnes). Our model can be readily extended to accommodate such constraints when operational-level planning is considered; refer to Appendix B for the detailed formulations.

3.2. Stochastic optimization model

We now propose a generalized stochastic model that more accurately reflects real-world uncertainty. In practice, certain decisions must be made before demand is realized, while others can be made based on the observed demand up to the decision point. In our problem, the signing of LTCs is a "here-and-now" decision: it must be finalized when the long-term supply plan is established without knowing the exact future demand. By contrast, spot market procurement decisions can be made adaptively in response to actual demand once it becomes known. Such decisions are often referred to as "wait-and-see" or "adjustable" decision variables. This distinction between here-and-now and wait-and-see decisions reflects the practical setting in which the decision- maker observes unfolding demand scenarios and adjusts short-term actions accordingly.

To incorporate these aspects into a stochastic framework, we introduce additional notation and parameters that account for demand uncertainty. This enables us to capture the multistage nature of the problem. Let $\tilde{d}^{jt}$ denote the random demand at terminal j in period t, for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$. Once the uncertainty is resolved in period t, the realized demand is denoted by d^{jt}. To keep track of the demand realization over time, we define $\tilde{\mathbf{d}}^t = (\tilde{d}^{jk}, \forall j \in \mathcal{J}, k \in \{1,\dots,t\})$ as the random vector of demands from periods 1 through t, and similarly let $\tilde{\mathbf{d}}^T$ denote the entire demand sequence over the planning horizon. For ease of exposition, we will use $\tilde{\mathbf{d}}$ to indicate $\tilde{\mathbf{d}}^T$. Once demands are realized, we write $\mathbf{d}^t = (d^{jk}, \forall j \in \mathcal{J}, k \in \{1,\dots,t\})$ for the sequence of observed demands through period t, and $\mathbf{d}$ for all realized demands from period 1 to T.

Among the decision variables $\boldsymbol{x}, \boldsymbol{\delta}, \boldsymbol{y}, \boldsymbol{s}, \boldsymbol{r}, \boldsymbol{w}, \boldsymbol{z}$ in $\mathrm{P_{DET}}$, all except $\boldsymbol{x}$ can be determined after observing realized demand portions, making them adjustable variables. In contrast, $\boldsymbol{x}$ must be fixed *before* the planning horizon because LTCs require advance commitments. This distinction arises because inventory, allocation, and spot market decisions can adapt dynamically to unfolding demand observations. The availability of demand information for each decision depends on its position within the decision timeline. For instance, inventory and allocation decisions s_{jt} and y_{jt} are made at the beginning of period t, and therefore depend only on demand observations up to period $t-1$. Accordingly, they are expressed as $s_{jt}(\tilde{\mathbf{d}}^{t-1})$ and $y_{jt}(\tilde{\mathbf{d}}^{t-1})$. Furthermore, the binary spot market activation decision δ_{jt} is made at the beginning of period t based on information up to period $t-1$, represented as $\delta_{jt}(\tilde{\mathbf{d}}^{t-1})$. Notably, $\delta_{jt}(\tilde{\mathbf{d}}^{t-1})$ is the only adjustable binary decision variable. In contrast, fulfillment decisions r_{jt}, procurement decisions from the spot market w_{jt}, and stockout quantities z_{jt} are determined later in period t, allowing them to incorporate the fully observed demand for that period. These are thus represented as $r_{jt}(\tilde{\mathbf{d}}^t)$, $w_{jt}(\tilde{\mathbf{d}}^t)$, and $z_{jt}(\tilde{\mathbf{d}}^t)$.

For ease of representation in the stochastic model, we define $\boldsymbol{\delta}(\tilde{\mathbf{d}}) = \left(\delta_{jt}(\tilde{\mathbf{d}}^{t-1}), \forall j \in \mathcal{J}, t \in \mathcal{T}\right)$ as the collection of adjustable binary decision variables. In addition, we introduce $\boldsymbol{\mu}(\tilde{\mathbf{d}})$ and $\boldsymbol{\nu}(\tilde{\mathbf{d}})$ to denote two distinct types of adjustable continuous decision variables respectively:

$$\boldsymbol{\mu}(\tilde{\mathbf{d}}) = \left(y_{jt}(\tilde{\mathbf{d}}^{t-1})\ \forall j \in \mathcal{J}, t \in \mathcal{T}, s_{jt}(\tilde{\mathbf{d}}^{t-1})\ \forall j \in \mathcal{J}, t \in \mathcal{T}^+\right),$$
$$\boldsymbol{\nu}(\tilde{\mathbf{d}}) = \left(r_{jt}(\tilde{\mathbf{d}}^t), w_{jt}(\tilde{\mathbf{d}}^t), z_{jt}(\tilde{\mathbf{d}}^t)\ \forall j \in \mathcal{J}, t \in \mathcal{T}\right).$$

With the above notations, the total cost of the LNG supply chain under a given demand $\mathbf{d}$ is defined as follows:

$$\boldsymbol{\Gamma}(\boldsymbol{x}, \boldsymbol{\delta}(\mathbf{d}), \boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d})) = \sum_{i\in\mathcal{I}} C_i x_i + \sum_{t\in\mathcal{T}} \left(\sum_{j\in\mathcal{J}} a_{jt} y_{jt}(\mathbf{d}^{t-1}) + \sum_{j\in\mathcal{J}} \left(f_{jt}\delta_{jt}(\mathbf{d}^{t-1}) + p_t w_{jt}(\mathbf{d}^t) + l_{jt} z_{jt}(\mathbf{d}^t) \right) + \sum_{j\in\mathcal{J}} h_t s_{jt+1}(\mathbf{d}^t) \right)$$

We propose a stochastic optimization model $\mathrm{P_{STOC}}$ that minimizes the expected total cost under demand uncertainty:

$(\mathrm{P_{STOC}})$

$$\min \quad \mathbb{E}_{\tilde{\mathbf{d}}} \left[\boldsymbol{\Gamma}\left(\boldsymbol{x}, \boldsymbol{\delta}(\tilde{\mathbf{d}}), \boldsymbol{\mu}(\tilde{\mathbf{d}}), \boldsymbol{\nu}(\tilde{\mathbf{d}})\right) \right] \tag{14}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{I}} U_{it} x_i = \sum_{j\in\mathcal{J}} y_{jt}(\tilde{\mathbf{d}}^{t-1}), \qquad \forall t \in \mathcal{T} \tag{15}$$

$$s_{jt}(\tilde{\mathbf{d}}^{t-1}) + y_{jt}(\tilde{\mathbf{d}}^{t-1}) \le \bar{s}_j, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{16}$$

$$\underline{s}_j \le s_{jt}(\tilde{\mathbf{d}}^{t-1}), \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}^+ \tag{17}$$

$$s_{jt+1}(\tilde{\mathbf{d}}^t) = (1-\beta_j)s_{jt}(\tilde{\mathbf{d}}^{t-1}) + y_{jt}(\tilde{\mathbf{d}}^{t-1}) - r_{jt}(\tilde{\mathbf{d}}^t), \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{18}$$

$$r_{jt}(\tilde{\mathbf{d}}^t) + w_{jt}(\tilde{\mathbf{d}}^t) + z_{jt}(\tilde{\mathbf{d}}^t) = \tilde{d}^{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{19}$$

$$w_{jt}(\tilde{\mathbf{d}}^t) \le M\delta_{jt}(\tilde{\mathbf{d}}^{t-1}), \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{20}$$

$$y_{jt}(\tilde{\mathbf{d}}^{t-1}) \ge 0, y_{jt}(\tilde{\mathbf{d}}^{t-1}) \in \mathcal{R}^{t-1}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{21}$$

$$w_{jt}(\tilde{\mathbf{d}}^t) \ge 0, w_{jt}(\tilde{\mathbf{d}}^t) \in \mathcal{R}^t, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{22}$$

$$r_{jt}(\tilde{\mathbf{d}}^t) \ge 0, r_{jt}(\tilde{\mathbf{d}}^t) \in \mathcal{R}^t, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{23}$$

$$z_{jt}(\tilde{\mathbf{d}}^t) \ge 0, z_{jt}(\tilde{\mathbf{d}}^t) \in \mathcal{R}^t, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{24}$$

$$x_i \in \{0,1\}, \qquad \forall i \in \mathcal{I}, \tag{25}$$

$$\delta_{jt}(\tilde{\mathbf{d}}^{t-1}) \in \mathcal{B}^{t-1}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}. \tag{26}$$

Here, $\mathcal{R}^t$ and $\mathcal{B}^t$ denote sets of functions that map $\mathbb{R}^{J\times t}$ to $\mathbb{R}$ and $\{0,1\}$, respectively, for each $t \in \mathcal{T}^+$. In addition, we assume $s_{j1}(\tilde{\mathbf{d}}^0) = s_{j1}$ is given for all $j \in \mathcal{J}$.

Note that $\mathrm{P_{STOC}}$ must be satisfied under every possible demand realization, thus requiring a solution that remains feasible across all scenarios. By solving $\mathrm{P_{STOC}}$, the decision maker obtains an optimal selection of LTCs and an adaptive strategy for allocation, fulfillment, and spot procurement that minimizes the expected total cost. However, $\mathrm{P_{STOC}}$ is inherently a multistage stochastic programming problem, which tends to be computationally intractable [45]. This intractability arises from two main factors. First, as the planning horizon T becomes large, the number of potential scenarios grows exponentially with T. Second, the presence of an adjustable binary decision variable $\boldsymbol{\delta}(\mathbf{d})$ further complicates the solution process, as the decision of whether to use the spot market may vary across different demand realizations.

Multistage stochastic optimization problems frequently arise in practical applications [46] and are widely recognized as critical, motivating the development of specialized solution approaches [47]. ARO has advanced over the years and demonstrates computational tractability in multistage inventory management problems [15]. However, incorporating adjustable binary variables within the ARO framework remains a major challenge, and several methodologies have been proposed to address this issue [18]. More recently, Lim et al. [33] introduced the TPA, which effectively handles adjustable binary decisions and exhibits strong computational performance in long-horizon problems. In light of these insights, we propose an ARO framework based on the TPA to address the LNG supply chain planning problem. This framework aims to ensure computational tractability and robust performance in handling adjustable binary spot market decisions over an extended multistage planning horizon.

4. Solution methodology

We solve $\mathrm{P_{STOC}}$ using a TPA based on RO. In Phase 1, we determine the binary decisions, namely the LTC selection $\boldsymbol{x}$ and the spot market usage $\boldsymbol{\delta}(\mathbf{d})$. This phase employs a TRO approach, which is designed to select binary decisions that accommodate as much demand uncertainty as possible while ensuring that the total cost remains below a predefined target. Once the binary decisions are fixed, Phase 2 optimizes the remaining adjustable continuous variables—namely allocation, fulfillment, and spot market procurement quantities—using an LDR. The LDR is a widely used strategy in ARO, enabling these decisions to adjust linearly based on the realized demand within the uncertainty set.

A key advantage of RO is that it does not require full knowledge of the probability distribution of uncertain demands. This property is particularly valuable in our setting, where the true distribution of future demand is unknown. Instead, we assume that the future demand lies within a predefined interval centered around a nominal forecast. By specifying this interval, we construct an uncertainty set that captures possible deviations from the forecast. In our strategic planning context, where decisions span a long horizon with bi-monthly periods, demand forecasts serve as the primary planning input, and the forecast error naturally defines the uncertainty interval through θ^{jt}. This interval formulation admits an exact reformulation of the robust counterpart as a tractable linear program, which is essential for the national-scale problem considered in this study. We note that the interval uncertainty set does not impose any distributional assumption on realized demand; Section 5.3 demonstrates that the framework remains effective across diverse distributional patterns, including asymmetric scenarios.

Formally, let each random demand $\tilde{d}^{jt}$ for terminal $j \in \mathcal{J}$ and period $t \in \mathcal{T}$ belong to a support set $[\underline{d}^{jt}, \overline{d}^{jt}]$ centered at the forecast value $\hat{d}^{jt}$. The lower and upper bounds are denoted as $\underline{d}^{jt} = \hat{d}^{jt} - b^{jt}$ and $\overline{d}^{jt} = \hat{d}^{jt} + b^{jt}$ where $b^{jt} = \theta^{jt}\hat{d}^{jt}$ for some $\theta^{jt} \in [0,1]$. Here, θ^{jt} represents the demand uncertainty parameter at terminal j in period t, capturing heterogeneous forecast accuracy across terminals and time periods. A larger θ^{jt} reflects lower forecast accuracy and widens the uncertainty interval, leading the model to adopt more conservative procurement decisions that hedge against a broader range of demand deviations at the expense of a higher expected total cost. Conversely, a smaller θ^{jt} reflects higher forecast confidence and narrows the interval, allowing the model to reduce the expected total cost while guaranteeing feasibility only against smaller demand deviations. Since the RO framework ensures feasibility for every demand realization within the interval, θ^{jt} directly governs the trade-off between the expected total cost and the degree of protection against forecast errors. Based on this construction, the uncertainty set for each $\tilde{d}^{jt}$ is given as follows:

$$D^{jt} = \left\{ d^{jt} \mid \hat{d}^{jt}(1-\theta^{jt}) \le d^{jt} \le \hat{d}^{jt}(1+\theta^{jt}) \right\}.$$

To describe the collection of all these sets up to period t, we define $\mathbf{D}^t := \left(D^{jk}, j \in \mathcal{J}, k \in \{1,\dots,t\}\right)$, and let $\mathbf{D} = \mathbf{D}^T$ denote the full uncertainty set over the entire planning horizon.

4.1. Phase 1 of TPA: Determining binary decisions

In Phase 1 of TPA, we utilize TRO to decouple the binary variables from $\mathrm{P_{STOC}}$. To proceed with the formulation, we define an adjustable uncertainty set that is parameterized by a scaling factor $\gamma \in [0,1]$, based on the original uncertainty set. For notational convenience, we define $D^{jt}(\gamma) := \left\{ d^{jt} \mid \hat{d}^{jt} - \gamma b^{jt} \le d^{jt} \le \hat{d}^{jt} + \gamma b^{jt} \right\}$. We then define the corresponding cumulative uncertainty set up to period t as $\mathbf{D}^t(\gamma) = \left(D^{jk}(\gamma), \forall j \in \mathcal{J}, k \in \{1,\dots t\}\right)$ and $\mathbf{D}^T(\gamma) = \mathbf{D}(\gamma)$.

We introduce a cost target parameter ξ as the overall budget for the LNG supply chain planning problem. A systematic methodology for determining an appropriate value of ξ will be proposed later in this work. Using this parameter, we derive the following TRO model ($\mathrm{P_{TRO}}$) by extending the stochastic programming model $\mathrm{P_{STOC}}$.

$$
\begin{aligned}
&(\mathrm{P_{TRO}}) \\
&\quad \gamma^* = \max\ \gamma \\
&\quad \text{s.t.}\ \ \boldsymbol{\Gamma}(\boldsymbol{x}, \boldsymbol{\delta}(\mathbf{d}), \boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d})) \le \xi, \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma) \\
&\qquad\quad \text{Constraints } (15)-(26), \qquad \forall \mathbf{d}^t \in \mathbf{D}^t(\gamma) \\
&\qquad\quad 0 \le \gamma \le 1
\end{aligned}
$$

In $\mathrm{P_{TRO}}$, we maximize γ to determine the largest adjustable uncertainty set $\mathbf{D}(\gamma)$ under which the solution remains robust. The cost function $\boldsymbol{\Gamma}$ quantifies the total cost resulting from binary decisions $\boldsymbol{x}, \boldsymbol{\delta}(\mathbf{d})$ and the adjustable continuous decisions $\boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d})$. By enforcing $\boldsymbol{\Gamma} \le \xi$, the model ensures that total costs remain within the predefined budget ξ for any demand realization in $\mathbf{D}(\gamma)$. Furthermore, the model imposes the feasibility of Constraints (15)–(24),(25)–(26) across all scenarios in the adjustable uncertainty set.

In the context of RO, a *static rule* refers to a decision policy where variables are fixed to constant values regardless of the realized uncertainty, i.e., $\boldsymbol{\pi}(\mathbf{d})=\boldsymbol{\pi}$. In contrast, an *adaptive rule* allows decisions to depend on observed uncertain parameters, i.e., $\boldsymbol{\pi}(\mathbf{d})$. Although adaptive rules offer greater flexibility, static rules provide computational tractability and can achieve optimality under specific structural conditions [21,26].

While the TRO model may also appear intractable, as it requires guaranteeing robust feasibility across all demand scenarios, Lim and Wang's theorem [25] demonstrates that a static rule $\boldsymbol{\delta}(\mathbf{d}) = \boldsymbol{\delta}$ is optimal if the worst-case scenario in the uncertainty set can be identified. However, this result cannot be directly applied to $\mathrm{P_{TRO}}$, as the presence of the non-adjustable variable $\boldsymbol{x}$ and Constraint (19) introduces structural complexities not addressed in their framework. Consequently, we modify the TPA method to accommodate these features.

Consider Constraint (19):

$$r_{jt}(\mathbf{d}^t) + w_{jt}(\mathbf{d}^t) + z_{jt}(\mathbf{d}^t) = \tilde{d}^{jt}.$$

Under a static rule, the left-hand side becomes a constant (independent of $\mathbf{d}$), while the right-hand side is a random variable. This mismatch makes a purely static rule infeasible, as it cannot satisfy all possible demand realizations exactly. To address this issue, we relax the TRO model $\mathrm{P}_{\mathrm{TRO}}$ as follows:

$$
\begin{aligned}
&(\mathrm{P}_{\mathrm{TRO-R}})\\
&\quad \gamma' = \max\ \gamma\\
&\quad \text{s.t.}\quad \boldsymbol{\Gamma}(\boldsymbol{x}, \boldsymbol{\delta}(\mathbf{d}), \boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d})) \le \xi, \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)\\
&\qquad\quad r_{jt}(\mathbf{d}^t) + w_{jt}(\mathbf{d}^t) + z_{jt}(\mathbf{d}^t) \ge d^{jt}, \qquad \forall \mathbf{d}^t \in \mathbf{D}^t(\gamma)\\
&\qquad\quad \text{Constraints (15)–(18), (20)–(26)}, \qquad \forall \mathbf{d}^t \in \mathbf{D}^t(\gamma)\\
&\qquad\quad 0 \le \gamma \le 1.
\end{aligned}
$$

Note that the relaxation of the second constraint leads to $\gamma' \ge \gamma^*$.

To generalize $\mathrm{P}_{\mathrm{TRO-R}}$, let $\boldsymbol{\pi}(\mathbf{d}) = \big(\boldsymbol{\delta}(\mathbf{d}), \boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d})\big)$ denote the set of all adjustable variables, and suppose $\boldsymbol{\pi}(\mathbf{d})$ lies within some feasible region $\boldsymbol{\Pi}$. This allows us to reformulate $\mathrm{P}_{\mathrm{TRO-R}}$ in its generalized form:

$$
\begin{aligned}
&(\mathrm{P}_{\mathrm{TRO-G}})\\
&\quad \gamma' = \max\ \gamma\\
&\quad \text{s.t.}\quad \boldsymbol{F}(\mathbf{d})\boldsymbol{\pi}(\mathbf{d}) \le \boldsymbol{g}(\mathbf{d}), \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)\\
&\qquad\quad \boldsymbol{H}(\mathbf{d})\{\boldsymbol{x}+\boldsymbol{\pi}(\mathbf{d})\} = \boldsymbol{k}(\mathbf{d}), \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)\\
&\qquad\quad \boldsymbol{\pi}(\mathbf{d}) \in \boldsymbol{\Pi}, \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)
\end{aligned}
$$

where $\boldsymbol{F}(\mathbf{d})$, $\boldsymbol{g}(\mathbf{d})$, $\boldsymbol{H}(\mathbf{d})$, and $\boldsymbol{k}(\mathbf{d})$ collectively represent deterministic and uncertain coefficients in $\mathrm{P}_{\mathrm{TRO-R}}$. Because $\mathrm{P}_{\mathrm{TRO-R}}$ is a relaxation of $\mathrm{P}_{\mathrm{TRO}}$, the continuous decisions obtained from this relaxed formulation are not feasible for the original problem. However, the binary decisions $(\boldsymbol{x}, \boldsymbol{\delta})$ for LTC selection and spot market usage remain valid regardless of the relaxation. In our two-phase approach, we utilize only these binary decisions from Phase 1 and determine the continuous decisions separately in Phase 2 using the LDR, which ensures feasibility for the original problem.

It is important to note that the second constraint in $\mathrm{P}_{\mathrm{TRO-G}}$ structurally diverges from the theorem proposed by Lim et al. [33], necessitating methodological refinements to preserve feasibility. To address this, we enforce a static rule by fixing $\boldsymbol{\pi}(\mathbf{d})=\boldsymbol{\pi}$ for all adjustable variables. This simplification reduces the problem to a computationally tractable form. The modified formulation is as follows:

$$
\begin{aligned}
&(\mathrm{P}_{\mathrm{TRO-S}})\\
&\quad \gamma^s = \max\ \gamma\\
&\quad \text{s.t.}\quad \boldsymbol{F}(\mathbf{d})\boldsymbol{\pi} \le \boldsymbol{g}(\mathbf{d}), \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)\\
&\qquad\quad \boldsymbol{H}(\mathbf{d})\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\mathbf{d}), \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma)\\
&\qquad\quad \boldsymbol{\pi} \in \boldsymbol{\Pi}, \qquad \forall \mathbf{d} \in \mathbf{D}(\gamma).
\end{aligned}
$$

However, the solution obtained from the static rule $(\boldsymbol{x},\boldsymbol{\pi})$ does not guarantee feasibility for all possible demand realizations within the uncertainty set $\mathbf{D}(\gamma)$. To address this limitation, we ensure the feasibility of $\mathrm{P}_{\mathrm{TRO-S}}$ by identifying the "worst-case scenario of uncertainty."

Definition 1 (Worst-case scenario of uncertainty (WSU)). Given a parameter γ and coefficients $\boldsymbol{F}(\mathbf{d}), \boldsymbol{g}(\mathbf{d}), \boldsymbol{H}(\mathbf{d}), \boldsymbol{k}(\mathbf{d})$, a special scenario $\check{\mathbf{d}}(\gamma) \in \mathbf{D}(\gamma)$ is called WSU if for each $(\boldsymbol{x},\boldsymbol{\pi}) \in \boldsymbol{\Pi}$ which satisfies $\boldsymbol{F}(\check{\mathbf{d}}(\gamma))\boldsymbol{\pi} \le \boldsymbol{g}(\check{\mathbf{d}}(\gamma))$ and $\boldsymbol{H}(\check{\mathbf{d}}(\gamma))\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\check{\mathbf{d}}(\gamma))$, it also satisfies $\boldsymbol{F}(\mathbf{d})\boldsymbol{\pi} \le \boldsymbol{g}(\mathbf{d})$ and $\boldsymbol{H}(\mathbf{d})\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\mathbf{d})$ for any $\mathbf{d} \in \mathbf{D}(\gamma)$.

In $\mathrm{P}_{\mathrm{TRO-R}}$, the WSU can be explicitly characterized. Given a parameter γ, we construct a collection of upper bound demand realizations $\hat{d}^{jt} + \gamma b^{jt}$ for all elements in the uncertainty set $\mathbf{D}(\gamma)$. This collection is denoted by $\check{\mathbf{d}}(\gamma) = (\hat{d}^{jt} + \gamma b^{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T})$ and we show that $\check{\mathbf{d}}(\gamma)$ represents the WSU for $\mathrm{P}_{\mathrm{TRO-R}}$. Suppose that the static rule $(\boldsymbol{x},\boldsymbol{\pi})$ satisfies all constraints in $\mathrm{P}_{\mathrm{TRO-R}}$ under this scenario. Then, we have $r_{jt} + w_{jt} \ge \hat{d}^{jt} + \gamma b^{jt} \ge d^{jt}$ for $\forall\, d^{jt} \in \mathbf{D}^t, \forall j \in \mathcal{J}, t \in \mathcal{T}$. This implies that the solution $(\boldsymbol{x},\boldsymbol{\pi})$ ensures feasibility for any demand realization within the uncertainty set $\mathbf{D}(\gamma)$. Furthermore, because the remaining constraints in $\mathrm{P}_{\mathrm{TRO-R}}$ are independent of specific demand realizations, the static rule $(\boldsymbol{x},\boldsymbol{\pi})$ remains feasible for all $\mathbf{d} \in \mathbf{D}(\gamma)$.

After identifying the WSU, we establish that the static rule $(\boldsymbol{x},\boldsymbol{\pi})$ serves as the optimal solution to $\mathrm{P}_{\mathrm{TRO-G}}$ through the following theorem.

Theorem 1 (Optimality of Static Rule). *For any $\gamma \in [0,1]$, if $\mathrm{P}_{\mathrm{TRO-G}}$ has a WSU denoted by $\check{\mathbf{d}}(\gamma)$, then the optimal solution $(\boldsymbol{x}^\dagger, \boldsymbol{\pi}^\dagger)$ for $\mathrm{P}_{\mathrm{TRO-G}}$ can be obtained by solving the following deterministic MILP model $\mathrm{P}_{\mathrm{TRO-T}}$:*

$$
\begin{aligned}
&(\mathrm{P}_{\mathrm{TRO-T}})\\
&\quad \gamma^\dagger = \max\ \gamma\\
&\quad \text{s.t.}\quad \boldsymbol{F}(\check{\mathbf{d}}(\gamma))\boldsymbol{\pi} \le \boldsymbol{g}(\check{\mathbf{d}}(\gamma)),\\
&\qquad\quad \boldsymbol{H}(\check{\mathbf{d}}(\gamma))\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\check{\mathbf{d}}(\gamma)),\\
&\qquad\quad \boldsymbol{\pi} \in \boldsymbol{\Pi}.
\end{aligned}
$$

Furthermore, the obtained $\gamma^\dagger$ satisfies the equality:

$$\gamma' = \gamma^s = \gamma^\dagger.$$

Proof. We begin by establishing the relationship among γ^s, $\gamma^\dagger$, and γ'. It is evident that $\gamma^s \le \gamma'$ because imposing a static rule is inherently more restrictive than allowing fully adaptive decision policies. However, we also show that $\gamma^s \ge \gamma^\dagger \ge \gamma'$, which together implies that all three values are equal.

From the definition of γ^s, we have:

$$
\begin{aligned}
\gamma^s &= \max\{\gamma \mid \boldsymbol{F}(\mathbf{d})\boldsymbol{\pi} \le \boldsymbol{g}(\mathbf{d}),\ \boldsymbol{H}(\mathbf{d})\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\mathbf{d}), \forall \mathbf{d} \in \mathbf{D}(\gamma),\ \boldsymbol{\pi} \in \boldsymbol{\Pi}\}\\
&\ge \max\left\{\gamma \mid \boldsymbol{F}(\check{\mathbf{d}}(\gamma))\boldsymbol{\pi} \le \boldsymbol{g}(\check{\mathbf{d}}(\gamma)),\ \boldsymbol{H}(\check{\mathbf{d}}(\gamma))\{\boldsymbol{x}+\boldsymbol{\pi}\} = \boldsymbol{k}(\check{\mathbf{d}}(\gamma)),\ \boldsymbol{\pi} \in \boldsymbol{\Pi}\right\} = \gamma^+.\\
&= \max\left\{\gamma \mid \boldsymbol{F}(\check{\mathbf{d}}(\gamma))\boldsymbol{\pi}(\check{\mathbf{d}}(\gamma)) \le \boldsymbol{g}(\check{\mathbf{d}}(\gamma)),\ \boldsymbol{H}(\check{\mathbf{d}}(\gamma))\{\boldsymbol{x}+\boldsymbol{\pi}(\check{\mathbf{d}}(\gamma))\} = \boldsymbol{k}(\check{\mathbf{d}}(\gamma))\right\}\\
&\ge \max\{\gamma \mid \boldsymbol{F}(\mathbf{d})\boldsymbol{\pi}(\mathbf{d}) \le \boldsymbol{g}(\mathbf{d}),\ \boldsymbol{H}(\mathbf{d})\{\boldsymbol{x}+\boldsymbol{\pi}(\mathbf{d})\} = \boldsymbol{k}(\mathbf{d}),\ \boldsymbol{\pi}(\mathbf{d}) \in \Pi, \forall \mathbf{d} \in \mathbf{D}(\gamma)\}\\
&= \gamma'.
\end{aligned}
$$

The first inequality follows from the definition of WSU. The last inequality holds because $\check{\mathbf{d}}(\gamma) \in \mathbf{D}(\gamma)$. Because we have established that $\gamma^s \le \gamma'$ and $\gamma^s \ge \gamma^\dagger \ge \gamma'$, it follows that: $\gamma^s = \gamma^\dagger = \gamma'$. This completes the proof, demonstrating that the optimal solution $(\boldsymbol{x}^\dagger, \boldsymbol{\pi}^\dagger)$ obtained from $\mathrm{P}_{\mathrm{TRO-T}}$ is also optimal for $\mathrm{P}_{\mathrm{TRO-G}}$, thereby confirming the optimality of the static rule. □

By Theorem 1, the optimal solution to $\mathrm{P}_{\mathrm{TRO-R}}$ can be obtained by solving the following deterministic MILP model:

$$
\begin{aligned}
&(\mathrm{P}_{\mathrm{STATIC}})\\
&\quad \gamma^\dagger = \max\quad \gamma\\
&\quad \text{s.t.}\quad \sum_{i\in\mathcal{I}} C_i x_i + \sum_{t\in\mathcal{T}}\left(\sum_{j\in\mathcal{J}} a_{jt} y_{jt} + \sum_{j\in\mathcal{J}} (f_{jt}\delta_{jt} + p_t w_{jt}) + \sum_{j\in\mathcal{J}} h_t s_{jt+1}\right) \le \xi\\
&\qquad\quad r_{jt} + w_{jt} + z_{jt} \ge \hat{d}^{jt} + \gamma b^{jt}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}\\
&\qquad\quad \text{Constraints (2)–(5), (7)–(13)},\\
&\qquad\quad 0 \le \gamma \le 1.
\end{aligned}
$$

By adjusting the cost target parameter ξ, the decision maker can control the level of conservativeness in the optimization model. A lower ξ leads to a more cost-efficient but potentially riskier decision, while a higher ξ results in a more conservative approach by allocating additional resources to hedge against uncertainty.

We note that an excessively small ξ (close to $\rho(0)$) may render the problem infeasible or yield a solution with weak robustness, as the resulting $\gamma^{\dagger}$ becomes too small to cover meaningful demand deviations. Conversely, an overly large ξ (close to $\rho(1)$) produces an unnecessarily conservative procurement plan that over-allocates resources against worst-case demand.

To determine an appropriate value for the cost target parameter ξ, we select a value within the range $[\rho(0), \rho(1)]$, where $\rho(\gamma)$ is computed by solving the following deterministic optimization problem:

$(\mathrm{P}_{\rho(\gamma)})$

$$\begin{aligned}\rho(\gamma) = \min\ & \sum_{i\in\mathcal{I}} C_i x_i + \sum_{t\in\mathcal{T}}\left(\sum_{j\in\mathcal{J}} a_{jt} y_{jt} + \sum_{j\in\mathcal{J}}(f_{jt}\delta_{jt} + p_t w_{jt}) + \sum_{j\in\mathcal{J}} h_t s_{jt+1}\right)\\ \text{s.t.}\ & r_{jt} + w_{jt} + z_{jt} \ge \hat{d}^{jt} + \gamma b^{jt}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}\\ & \text{Constraints (2)–(5), (7)–(13),}\\ & 0 \le \gamma \le 1.\end{aligned}$$

To flexibly control the level of conservativeness, we introduce a parameter $\alpha \in [0,1]$, which is called the target coefficient, and compute ξ as a convex combination of the two extreme values: $\xi = \alpha\rho(0) + (1-\alpha)\rho(1)$. When α is close to 1, the budget ξ approaches $\rho(0)$ and the resulting $\gamma^{\dagger}$ is small, so the solution covers only a narrow band of demand uncertainty. As α decreases toward 0, ξ increases toward $\rho(1)$, allowing $\gamma^{\dagger}$ to grow and the model to hedge against a wider range of demand realizations. Table S1 in the Supplementary Material empirically demonstrates this relationship between α and $\gamma^{\dagger}$.

Once the appropriate value of ξ is determined, we proceed by solving the static optimization problem $\mathrm{P}_{\mathrm{STATIC}}$. This yields the optimal binary decisions $(\boldsymbol{x}^*, \boldsymbol{\delta}^*)$, which define key strategic choices such as LTC selection and spot market usage. These binary decisions are passed to Phase 2, where the remaining adjustable variables are determined using an LDR.

4.2. Phase 2 of TPA: Determining continuous decisions

In Phase 2 of TPA, we utilize the fixed binary decisions $(\boldsymbol{x}^*, \boldsymbol{\delta}^*)$ obtained from Phase 1 to determine the adjustable continuous decisions. Building upon these strategic choices, we develop a policy that adaptively determines allocation, fulfillment, and spot market procurement as demand uncertainty unfolds over time. Specifically, the adaptive nature of the policy enables it to leverage updated demand information in each period, thereby supporting more responsive and accurate operational decisions. This decision process is formulated as the following RO model, which minimizes the worst-case expected total cost under the predefined uncertainty sets $\mathbf{D}$:

$(\mathrm{P}_{\mathrm{ARO}})$

$$\begin{aligned}\min\max_{\mathbf{d}\in\mathbf{D}}\ & \mathbb{E}_{\mathbf{d}}\Bigg[\sum_{i\in\mathcal{I}} C_i x_i^* + \sum_{t\in\mathcal{T}}\Bigg(\sum_{j\in\mathcal{J}} a_{jt} y_{jt}(\mathbf{d}^{t-1}) + \sum_{j\in\mathcal{J}}\Big(f_{jt}\delta_{jt}^* + p_t w_{jt}(\mathbf{d}^t)\\ & \qquad + l_{jt} z_{jt}(\mathbf{d}^t)\Big) + \sum_{j\in\mathcal{J}} h_t s_{jt+1}(\mathbf{d}^{t-1})\Bigg)\Bigg]\\ \text{s.t.}\ & \sum_{i\in\mathcal{I}} U_{it} x_i^* = \sum_{j\in\mathcal{J}} y_{jt}(\mathbf{d}^{t-1}), \qquad \forall t\in\mathcal{T},\ \forall \mathbf{d}^{t-1}\in\mathbf{D}^{t-1}\\ & w_{jt}(\mathbf{d}^t) \le M\delta_{jt}^*, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T},\ \forall \mathbf{d}^t\in\mathbf{D}^t\\ & \text{Constraints (16)–(19), (21)–(24),} \qquad \forall \mathbf{d}^t\in\mathbf{D}^t.\end{aligned}$$

However, solving $\mathrm{P}_{\mathrm{ARO}}$ in its full generality is typically intractable due to the high dimensionality and complexity introduced by adjustable decisions under uncertainty. To address this challenge, we adopt an LDR, which restricts the functional form of the adjustable continuous decision variables while incurring only a minimal loss in optimality. It is important to note that the LDR based reformulation yields an approximation of the original ARO model (i.e., $\mathrm{P}_{\mathrm{ARO}}$), not a mathematical equivalence. A true optimal policy could be any measurable function of realized uncertainty. However, by restricting adjustable decisions to affine functions, we transform the intractable functional optimization into a tractable LP model. Although this restriction may exclude nonlinear optimal policies, See and Sim [21] demonstrate that affine policies achieve near-optimal performance in multi-period inventory management contexts. Specifically, each adjustable variable is represented as an affine function of the realized demand up to the corresponding decision point, as illustrated below:

$$y_{jt}\left(\mathbf{d}^{t-1}\right) = y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k} d^{j'k}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}$$

$$s_{jt}\left(\mathbf{d}^{t-1}\right) = s_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} s_{jt}^{j'k} d^{j'k}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}^+$$

$$r_{jt}\left(\mathbf{d}^{t}\right) = r_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} r_{jt}^{j'k} d^{j'k}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}$$

$$w_{jt}\left(\mathbf{d}^{t}\right) = w_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} w_{jt}^{j'k} d^{j'k}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}$$

$$z_{jt}\left(\mathbf{d}^{t}\right) = z_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} z_{jt}^{j'k} d^{j'k}, \qquad \forall j\in\mathcal{J}, t\in\mathcal{T}$$

Note that $s_{jt}\left(\mathbf{d}^{t-1}\right)$ serves as a dependent variable and can likewise be represented as an affine function of the realized demands. Once the coefficients $\left\{y_{jt}^0, y_{jt}^{j'k}, s_{jt}^0, s_{jt}^{j'k}, r_{jt}^0, r_{jt}^{j'k}, w_{jt}^0, w_{jt}^{j'k}, z_{jt}^0, z_{jt}^{j'k}\right\}$ are specified, the corresponding demand-adjustable decisions are determined dynamically based on the observed demands up to the decision point. Thus, rather than determining the operational decisions explicitly in each scenario, we optimize over the set of possible coefficient values that define the decision rules. In other words, we are designing a *policy*, not individual decisions, ensuring that the system adapts in a linear manner as uncertainty unfolds. The coefficients for LDRs introduced in this section are summarized in Appendix A.

Specifically, since LDRs are affine in demand, each decision variable evaluated at the nominal demand $\hat{d}^{j'k}$ equals its expectation. For instance, for the allocation decision:

$$\mathbb{E}_{\mathbf{d}}[y_{jt}(\mathbf{d}^{t-1})] = y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k}\mathbb{E}_{\mathbf{d}}[d^{j'k}] = y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k}\hat{d}^{j'k}.$$

Combined with the linear cost structure, the expected total cost reduces to a deterministic linear function of the LDR coefficients and nominal demands $\hat{d}^{j'k}$. This eliminates the worst-case maximization over $\mathbf{d}\in\mathbf{D}$ in the objective. Instead, robustness is enforced through constraints that must hold for all demand realizations, as formulated in $\mathrm{P}_{\mathrm{LDR}}$ below.

We note that this tractable reformulation relies on the linear cost structure, which is well-supported at the strategic planning level [11, 36]. Nonlinear cost effects, such as terminal congestion or economies of scale, are more relevant at operational time scales and would require solution approaches beyond LDRs (refer to Section 6). By applying this LDR framework to $\mathrm{P}_{\mathrm{ARO}}$, we reformulate the problem as $\mathrm{P}_{\mathrm{LDR}}$, where the objective is to select the coefficient values that minimize the total cost while ensuring feasibility for all demand realizations.

($\mathrm{P_{LDR}}$)

$$\min \sum_{i\in\mathcal{I}} C_i x_i^* + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} a_{jt}\left(y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k}\hat{d}^{j'k}\right) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}}\left(f_{jt}\delta_{jt}^* + p_t\left(w_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} w_{jt}^{j'k}\hat{d}^{j'k}\right)\right) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} h_t\left(s_{jt+1}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} s_{jt+1}^{j'k}\hat{d}^{j'k}\right)\Bigg) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} l_{jt}\left(z_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} z_{jt}^{j'k}\hat{d}^{j'k}\right)\Bigg)$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{I}} U_{it}x_i = \sum_{j\in\mathcal{J}} y_{jt}^0, \quad \forall t\in\mathcal{T} \tag{27}$$

$$\sum_{j\in\mathcal{J}} y_{jt}^{j'k} = 0, \quad \forall j'\in\mathcal{J},\, t\in\mathcal{T},\, k=1,\dots,t-1 \tag{28}$$

$$s_{jt}^0 + y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}(s_{jt}^{j'k} + y_{jt}^{j'k})d^{j'k} \le \bar{s}_j, \ \forall j'\in\mathcal{J},\, t\in\mathcal{T},\, \mathbf{d}\in\mathbf{D}^{t-1} \tag{29}$$

$$\underline{s}_j \le s_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} s_{jt}^{j'k}d^{j'k}, \quad \forall j'\in\mathcal{J},\, t\in\mathcal{T}^+,\, \mathbf{d}\in\mathbf{D}^{t-1} \tag{30}$$

$$s_{jt+1}^0 = (1-\beta_j)s_{jt}^0 + y_{jt}^0 - r_{jt}^0, \quad \forall j\in\mathcal{J},\, t\in\mathcal{T} \tag{31}$$

$$s_{jt+1}^{j'k} = (1-\beta_j)s_{jt}^{j'k} + y_{jt}^{j'k} - r_{jt}^{j'k}, \quad j\in\mathcal{J},\, j'\in\mathcal{J},\, t\in\mathcal{T},\, k=1,\dots,t-1 \tag{32}$$

$$s_{jt+1}^{j't} + r_{jt}^{j't} = 0, \quad \forall j\in\mathcal{J},\, j'\in\mathcal{J},\, t\in\mathcal{T} \tag{33}$$

$$r_{jt}^0 + w_{jt}^0 + z_{jt}^0 = 0, \quad \forall j\in\mathcal{J},\, t\in\mathcal{T} \tag{34}$$

$$r_{jt}^{j'k} + w_{jt}^{j'k} + z_{jt}^{j'k} = 0, \quad j\in\mathcal{J},\, j'\in\mathcal{J},\, t\in\mathcal{T},\, k=1,\dots,t-1, \tag{35}$$

$$r_{jt}^{jt} + w_{jt}^{jt} + z_{jt}^{jt} = 1, \quad \forall j\in\mathcal{J},\, t\in\mathcal{T}, \tag{36}$$

$$w_{jt}^0 = 0, \quad (j,t)\in\{(j,t)\mid \delta_{jt}^* = 0\}, \tag{37}$$

$$w_{jt}^{j'k} = 0, \quad (j,t)\in\{(j,t)\mid \delta_{jt}^* = 0\},\, j'\in\mathcal{J},\, k=1,\dots,t \tag{38}$$

$$0 \le w_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} w_{jt}^{j'k}d^{j'k} \le M, \quad (j,t)\in\{(j,t)\mid \delta_{jt}^* = 1\},\, \mathbf{d}\in\mathbf{D}^t, \tag{39}$$

$$y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k}d^{j'k} \ge 0, \quad j\in\mathcal{J},\, t\in\mathcal{T},\, \mathbf{d}\in\mathbf{D}^{t-1}, \tag{40}$$

$$r_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} r_{jt}^{j'k}d^{j'k} \ge 0, \quad j\in\mathcal{J},\, t\in\mathcal{T},\, \mathbf{d}\in\mathbf{D}^t, \tag{41}$$

$$z_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} z_{jt}^{j'k}d^{j'k} \ge 0, \quad j\in\mathcal{J},\, t\in\mathcal{T},\, \mathbf{d}\in\mathbf{D}^t. \tag{42}$$

Note that Constraints (29)–(30) and (39)–(41) remain valid for all demand realizations under the affine decision framework. While these constraints may appear to increase the complexity of the RO problem, they can be systematically reformulated through the application of duality theory.

The stochastic constraints (29)–(30) and (39)–(41) can be written in the following general form:

$$\lambda_0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} \lambda^{j'k}d^{j'k} \le 0, \quad \forall\, \mathbf{d}\in\mathbf{D}^{t-1}.$$

Considering all possible demand realizations within the predefined uncertainty set, this constraint is equivalent to the following robust inequality:

$$\max_{\mathbf{d}}\left\{\lambda_0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}\lambda^{j'k}d^{j'k} \,:\, d^{j'k}\in[(1-\theta^{j'k})\hat{d}^{j'k}, (1+\theta^{j'k})\hat{d}^{j'k}]\right\} \le 0.$$

This robust constraint can be equivalently reformulated into a set of linear inequalities by introducing new nonnegative dual variables $\eta^{j'k} \ge 0$ for all $j'\in\mathcal{J}, k=1,\dots,t-1$:

$$\lambda_0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}\left(\lambda^{j'k}\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta^{j'k}\right) \le 0,$$

$$-\eta^{j'k} \le \lambda^{j'k} \le \eta^{j'k} \quad \forall j'\in\mathcal{J}, t\in\mathcal{T}.$$

Through this process, we ultimately reformulate $\mathrm{P_{LDR}}$ as the following LP model:

($\mathrm{P_{LDR-LP}}$)

$$\min \sum_{i\in\mathcal{I}} C_i x_i^* + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} a_{jt}(y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1} y_{jt}^{j'k}\hat{d}^{j'k}) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}}(f_{jt}\delta_{jt}^* + p_t(w_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} w_{jt}^{j'k}\hat{d}^{j'k})) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} h_t(s_{jt+1}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} s_{jt+1}^{j'k}\hat{d}^{j'k})) + \sum_{j\in\mathcal{J}}\sum_{t\in\mathcal{T}} l_{jt}(z_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t} z_{jt}^{j'k}\hat{d}^{j'k}))$$

$$\text{s.t.} \quad s_{jt}^0 + y_{jt}^0 - \bar{s}_j + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}\left[(s_{jt}^{j'k} + y_{jt}^{j'k})\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta_{1jt}^{j'k}\right] \le 0, \quad \forall j\in\mathcal{J}, t\in\mathcal{T} \tag{43}$$

$$-\eta_{1jt}^{j'k} \le s_{jt}^{j'k} + y_{jt}^{j'k} \le \eta_{1jt}^{j'k}, \quad \forall j,j'\in\mathcal{J}, t\in\mathcal{T}, k=1,\dots,t-1, \tag{44}$$

$$\underline{s}_j - s_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}\left[(-s_{jt}^{j'k})\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta_{2jt}^{j'k}\right] \le 0, \quad \forall j\in\mathcal{J}, t\in\mathcal{T}^+ \tag{45}$$

$$-\eta_{2jt}^{j'k} \le -s_{jt}^{j'k} \le \eta_{2jt}^{j'k}, \quad \forall j,j'\in\mathcal{J}, t\in\mathcal{T}, k=1,\dots,t-1, \tag{46}$$

$$-w_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t}\left[(-w_{jt}^{j'k})\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta_{3jt}^{j'k}\right] \le 0, \quad (j,t)\in\{(j,t)\mid\delta_{jt}^* = 1\}, \tag{47}$$

$$-\eta_{3jt}^{j'k} \le -w_{jt}^{j'k} \le \eta_{3jt}^{j'k}, \quad (j,t)\in\{(j,t)\mid\delta_{jt}^* = 1\}, \forall j'\in\mathcal{J}, k=1,\dots,t, \tag{48}$$

$$w_{jt}^0 - M + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t}\left[(w_{jt}^{j'k})\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta_{4jt}^{j'k}\right] \le 0, \quad (j,t)\in\{(j,t)\mid\delta_{jt}^* = 1\}, \tag{49}$$

$$-\eta_{4jt}^{j'k} \le w_{jt}^{j'k} \le \eta_{4jt}^{j'k}, \quad (j,t)\in\{(j,t)\mid\delta_{jt}^* = 1\}, \forall j'\in\mathcal{J}, k=1,\dots,t, \tag{50}$$

$$-y_{jt}^0 + \sum_{j'\in\mathcal{J}}\sum_{k=1}^{t-1}\left[(-y_{jt}^{j'k})\hat{d}^{j'k} + \theta^{j'k}\hat{d}^{j'k}\eta_{5jt}^{j'k}\right] \le 0, \quad \forall j\in\mathcal{J}, t\in\mathcal{T} \tag{51}$$

$$-\eta_{5jt}^{j'k} \le -y_{jt}^{j'k} \le \eta_{5jt}^{j'k}, \quad \forall j,j'\in\mathcal{J}, t\in\mathcal{T}, k=1,\dots,t-1, \tag{52}$$

$$-r_{jt}^{0}+\sum_{j'\in\mathcal{J}}\sum_{k=1}^{t}\left[(-r_{jt}^{j'k})\,\hat{d}^{j'k}+\theta^{j'k}\,\hat{d}^{j'k}\,\eta_{6jt}^{j'k}\right]\le 0,\quad \forall j\in\mathcal{J},t\in\mathcal{T} \tag{53}$$

$$-\eta_{6jt}^{j'k}\le -r_{jt}^{j'k}\le \eta_{6jt}^{j'k},\qquad \forall j,j'\in\mathcal{J},t\in\mathcal{T},k=1,\dots,t, \tag{54}$$

$$-z_{jt}^{0}+\sum_{j'\in\mathcal{J}}\sum_{k=1}^{t}\left[(-z_{jt}^{j'k})\,\hat{d}^{j'k}+\theta^{j'k}\,\hat{d}^{j'k}\,\eta_{7jt}^{j'k}\right]\le 0,\quad \forall j\in\mathcal{J},t\in\mathcal{T} \tag{55}$$

$$-\eta_{7jt}^{j'k}\le -z_{jt}^{j'k}\le \eta_{7jt}^{j'k},\qquad \forall j,j'\in\mathcal{J},t\in\mathcal{T},k=1,\dots,t, \tag{56}$$

$$\eta_{1jt}^{j'k},\eta_{2jt}^{j'k},\eta_{3jt}^{j'k},\eta_{4jt}^{j'k},\eta_{5jt}^{j'k},\eta_{6jt}^{j'k},\eta_{7jt}^{j'k}\ge 0,\quad \forall j,j'\in\mathcal{J},t\in\mathcal{T},k=1,\dots,t, \tag{57}$$

$$\text{Constraints (27)–(28), (31)–(38).} \tag{58}$$

Although $\mathrm{P_{STOC}}$ initially appears computationally intractable due to its multistage stochastic structure and the presence of adjustable binary variables, applying the TPA enables a tractable reformulation. Phase 1 solves a MILP to determine binary decisions, and Phase 2 solves a linear program for continuous decisions. As a result, the original complex stochastic problem can now be handled using standard optimization solvers, making it possible to support nationwide LNG supply chain planning with practical computational tools.

5. Computational experiments

In this section, we evaluate the performance and practical applicability of our proposed approach, the TPA model, through a series of computational experiments. While the model exhibits strong methodological properties, our primary focus is to assess its potential for real-world application in nationwide LNG supply chain planning. To this end, we explore the following primary research questions:

- Does the TPA model provide practical and cost-effective policies for nationwide LNG supply chain planning?
- How do supply chain decisions under the TPA model differ from those under forecast-based planning?
- How does the accuracy of demand forecasts affect LNG supply chain planning under the TPA model?
- How do key supply chain parameters affect procurement decisions under the TPA model?

In addition, we aim to derive managerial insights that can inform strategic decision-making in LNG supply chain planning. Sections 5.1 and 5.2 address Research Questions 1 and 2 by analyzing cost performance and supply chain decisions. Section 5.3 investigates Research Question 3 through forecast accuracy experiments, and Section 5.4 examines Research Question 4 via sensitivity analysis. Finally, Section 5.5 derives managerial insights from these findings.

We begin by outlining the configuration of model parameters. Parameter values were selected based on multiple data sources (e.g., [1,2,48–50]) to approximate real-world conditions. This study addresses LNG supply chain planning in South Korea using bi-monthly decision intervals. To rigorously assess computational performance, Section 5.1 employs a 10-year planning horizon ($T = 60$). For the subsequent detailed analyses on cost structure and parameters (Sections 5.2–5.4), we adopt a 5-year horizon ($T = 30$). Currently, Korea Gas Corporation—the state-owned enterprise responsible for managing South Korea's LNG supply chain—operates five major LNG terminals; accordingly, we set $J = 5$ to represent the number of terminals. Furthermore, we consider multiple potential LTCs as feasible supply options. To evaluate computational scalability, we vary the number of contract candidates (I) across experiments.

The supply terms of each LTC candidate were generated within realistic bounds to reflect typical industry practices. Each contract specifies a per-period supply quantity (SUP), its start and end periods (START, END), a price per tonne of LNG (PRICE), and a fixed cost (FC) covering the entire contract duration. Table 2 summarizes the parameter ranges.

Table 2
Parameter settings.

Parameter (Integer)	Value	Parameter (Real-valued)	Value
SUP	$[15, 20]\times 10^5$	a_{jt}	[20, 25]
START	$[1, \lfloor T/2\rfloor]$	h_t	[40, 60]
END	$[\text{START}+\lfloor T/2\rfloor, T]$	p_t	$[8, 10]\times 10^2$
PRICE	[585, 600]	f_{jt}	10^7
FC	$[10, 15]\times 10^7$	l_{jt}	$[16, 20]\times 10^3$

As noted earlier, LNG is susceptible to natural evaporation (BOG), and the inventory balance constraint (Constraint (5)) incorporates a BOG loss parameter $\beta_j \in [0,1]$ to capture the net loss rate after recovery at terminal j per period. Modern Korean terminals achieve BOG recovery rates of 95–99% through reliquefaction [51,52], resulting in effective losses below 1%. Accordingly, we set $\beta_j = 0$ as the baseline in our experiments while retaining the general formulation. Section 5.1 examines the impact of nonzero BOG loss rates across all problem sizes, confirming that this simplification does not affect the comparative findings.

The aggregate national LNG demand for South Korea in 2024, which amounts to 4.6×10^7 tonnes, is established as the total forecasted demand for year 1. Subsequently, we assume that the annual demand will increase by 0.2% each year over the remaining 4 years of the planning horizon. Each year is divided into six bi-monthly periods. Following standard industry practice, the seasonal demand distribution is characterized by a 3:2:1 ratio for winter (periods 1 and 6), spring/autumn (periods 2 and 5), and summer (periods 3 and 4), respectively. Accordingly, period 1, corresponding to the first bi-monthly period of Year 1, accounts for one-quarter of the annual demand in that specific year. For each period, the sum of demands across all terminals equals the total demand in that period. Each terminal accounts for a fixed proportion of the total demand for that period.

In practice, LNG spot prices often exhibit a positive correlation with aggregate demand. For instance, spot prices in Asian markets tend to spike during winter heating seasons when demand peaks. We can capture this co-movement by setting the spot prices correlated with the demand, e.g., $p_t = \varphi \sum_{j\in\mathcal{J}} \hat{d}^{jt}$, where φ is calibrated from historical spot prices (refer to Appendix B). However, in our computational experiments, we deliberately adopt independently sampled spot prices (Table 2) to provide a more rigorous test of algorithmic performance. For each terminal, the demand proportions and storage capacity are presented in Table 3. The maximum capacity $\bar{s}_j$ and minimum inventory level $\underline{s}_j$ are determined based on the storage capacity of each terminal by multiplying the respective multipliers.

To evaluate the proposed methodology, we conducted computational experiments under the following settings. The TPA model was implemented in Python 3.13.5 and solved using Gurobi Optimizer 12.0.3. For both Phase 1 and Phase 2, the optimization procedures were configured with a relative optimality gap tolerance of 1% and a time limit of 3600 seconds referring to previous studies [33,35]. All other solver parameters were kept at their default values.

In Phase 1 of the TPA, we conducted experiments with eleven different values of α to identify an appropriate cost budget ξ, varying the parameter α from 0.0 to 1.0 in increments of 0.1. In Phase 2, we

Table 3
Terminal-wise demand shares and capacity bounds.

Terminal	Demand share (%)	Storage capacity ($\times 10^3$ tonnes)
1	17	1408
2	20	1760
3	25	2288
4	25	2068
5	13	1232

determined the coefficients of LDRs for continuous variables based on the binary solutions obtained from Phase 1 ($\boldsymbol{x}$ and $\boldsymbol{\delta}$), thereby constructing a policy for each adjustable variable. Once actual demand is realized, these policies prescribe specific actions in each period. All experiments were executed on a PC equipped with an Intel(R) Xeon(R) w3-2435 (3.10 GHz), 256 GB RAM, and Windows 11 Pro for Workstations (64-bit).

5.1. Performance analysis

To evaluate the effectiveness of the TPA model, we compare its performance against three benchmark models: PI, DET, and SPOT. The PI model represents an idealized case with *perfect information*, assuming that all future demand realizations are known in advance. By solving P_{DET} under these perfect forecasts, the PI model achieves the minimum possible total cost and thus serves as a theoretical lower bound.

The SPOT model, on the other hand, operates entirely without LTCs and relies solely on reactive spot market procurement to meet demand. Because spot prices are generally higher than contracted prices, the SPOT model provides the highest total cost among the benchmarks, thus serving as a practical upper bound. Together with the PI model, it defines a cost range for evaluating the performance of the TPA model.

The DET model serves as a practical benchmark that replicates current industry practices (see Algorithm S2 in the Supplementary Material for detailed pseudocode). First, it solves P_{DET} once using nominal demand forecasts $\hat{\mathbf{d}}$ to determine the LTCs $\boldsymbol{x}^*$. After that, rather than re-optimizing for each realized demand scenario, the model simulates period-by-period decisions using a target inventory heuristic (Algorithm S5 in the Supplementary Material). Specifically, LNG arriving under LTCs is allocated proportionally to each terminal according to its share of total demand. The target inventory level is dynamically determined based on forecasted future demand using a lookahead strategy. If the remaining inventory falls below the target level after fulfilling demand, the shortfall is covered via spot procurement. Conversely, any surplus is carried over to the next period. This two-stage structure reflects real-world LNG supply management, where LTCs are fixed in advance based on forecasts, while operational decisions are made adaptively through heuristic rules in response to actual conditions.

For simulation, we generate demand realizations within the predefined uncertainty set. Although the exact probability distribution of demand is unknown, we sample demand deviations from a series of Beta distributions to produce diverse demand patterns within the allowable bounds. Specifically, for each scenario, we generate the realized demand d^{jt} as $d^{jt} = \hat{d}^{jt}(1 - \theta^{jt} + 2\theta^{jt}\epsilon)$, where $\epsilon \sim \text{Beta}(a, b)$ is a random variable defined on the interval $[0, 1]$. We adopt the Beta distribution for its parametric flexibility: Beta(1,1) equals the uniform distribution, Beta(2,2) approximates a symmetric bell-shaped pattern similar to the truncated normal, and Beta(0.5,2) and Beta(2,0.5) capture left-skewed and right-skewed demand patterns, respectively.

To evaluate model performance, we conduct Monte Carlo simulation experiments with 100 demand scenarios for each Beta distribution. For each scenario, all six models—PI, TPA, DET, SPOT, STATIC, and SAA—are evaluated under the same realized demand $\mathbf{d}^{(s)}$ where $s \in \{1, \dots, 100\}$ to ensure a fair comparison (see Algorithms S1–S5 in the Supplementary Material). The PI model solves P_{DET} for each scenario with perfect demand knowledge. The DET model uses pre-determined contracts from the contract selection stage and simulates operational decisions under realized demand. The TPA model applies derived LDRs. The SPOT model simulates spot-market-only operations. The STATIC model represents the worst-case cost from Phase 1 ($P_{\rho(\gamma)}$ with $\gamma = 1$), where all decisions are made before uncertainty realization, serving as a conservative static policy benchmark. The Sample Average Approximation (SAA) model represents a stochastic programming benchmark that optimizes over $N = 100$ sampled scenarios from an assumed demand distribution (see Algorithm S4 in the Supplementary Material). To ensure the robustness of the PI benchmark to sampling error, we conducted a convergence analysis with 1000 scenarios (see Figure S1 in the Supplementary Material for details). The coefficient of variation (CV) remains below 1% across all distributional shapes, confirming that our PI estimate provides a statistically reliable lower bound. We set $\theta^{jt} = \theta = 0.1$ for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$, and Table 4 reports the mean total cost and optimality gap across all 100 scenarios for each distribution.

Table 4 presents the total cost and optimality gap for each model under five different demand distributions. The PI model consistently yields the lowest total cost, serving as a theoretical lower bound, while the SPOT model results in the highest cost, reflecting the absence of LTCs. Between these extremes, the TPA model clearly dominates the industry-standard DET benchmark model. The TPA model achieves an optimality gap below 5% in most test cases, with only one exception under the asymmetric Beta(0.5, 2) distribution with $I = 20$. In contrast, the DET model exhibits significantly higher gaps ranging from 7% to 14%, and the STATIC model performs even more conservatively because all decisions are fixed before uncertainty realization. Among the benchmarks, SAA improves upon DET by optimizing over sampled

Table 4
Total cost and optimality gap results across models under different demand distributions.

I	Distribution	PI		TPA			DET		SPOT		STATIC		SAA		
		Cost	Gap	Cost	Gap	CPUs	Cost	Gap	Cost	Gap	Cost	Gap	Cost	Gap	CPUs
10	Beta(0.5, 0.5)	3.27	0.00	3.33	2.07	285.4	3.52	7.91	4.52	38.37	3.57	9.17	3.51	7.35	54.4
	Beta(1, 1)	3.26	0.00	3.33	2.06		3.51	7.46	4.50	38.02		9.51	3.50	7.12	64.3
	Beta(2, 2)	3.26	0.00	3.33	2.05		3.50	7.14	4.49	37.48		9.51	3.49	6.91	51.6
	Beta(0.5, 2)	3.03	0.00	3.10	2.20		3.28	8.00	4.21	38.91		17.82	3.26	7.34	26.5
	Beta(2, 0.5)	3.45	0.00	3.56	3.23		3.77	9.35	4.83	40.07		3.48	3.70	7.15	46.9
20	Beta(0.5, 0.5)	3.19	0.00	3.32	4.16	462.8	3.49	9.22	4.47	40.03	3.52	10.34	3.43	7.31	60.8
	Beta(1, 1)	3.19	0.00	3.33	4.22		3.45	8.07	4.45	39.46		10.34	3.41	6.74	61.7
	Beta(2, 2)	3.19	0.00	3.33	4.16		3.44	7.74	4.44	39.17		10.34	3.41	6.65	68.9
	Beta(0.5, 2)	2.97	0.00	3.16	6.34		3.19	7.66	4.17	40.40		18.52	3.18	7.23	37.2
	Beta(2, 0.5)	3.40	0.00	3.50	2.77		3.78	11.17	4.79	40.72		3.53	3.77	10.67	83.8
30	Beta(0.5, 0.5)	3.15	0.00	3.30	4.97	885.8	3.47	10.25	4.53	43.79	3.50	11.11	3.39	7.79	31.5
	Beta(1, 1)	3.15	0.00	3.30	4.92		3.45	9.41	4.51	43.17		11.11	3.38	7.26	21.8
	Beta(2, 2)	3.15	0.00	3.30	4.99		3.42	8.59	4.49	42.74		11.11	3.37	7.20	24.2
	Beta(0.5, 2)	2.98	0.00	3.12	4.57		3.41	14.61	4.21	41.40		17.45	3.18	6.79	57.6
	Beta(2, 0.5)	3.36	0.00	3.49	4.07		3.73	11.05	4.84	44.19		4.17	3.59	6.78	16.2

Note: - Cost values are in $\times 10^{11}$($); Gap values are in %; CPUs values are in seconds. - TPA uses optimal α: 0.8 for $I = 10$, 0.4 for $I = 20$, 0.5 for $I = 30$. - STATIC cost is identical across distributions for each I because it is computed before uncertainty realization (Phase 1 worst-case cost with $\gamma = 1$). - SAA must be solved separately for each distribution.

Table 5
Average optimality gap (%) under different BOG loss rates and problem sizes.

I	β_j	PI	Optimality gap (%)				
		Cost	TPA	DET	SPOT	STATIC	SAA
10	0.000	3.25	2.32	7.94	38.65	9.92	7.20
	0.015	3.28	2.08	9.50	37.90	9.79	8.93
	0.030	3.30	1.50	12.60	38.82	9.51	10.30
20	0.000	3.19	4.32	8.78	39.96	10.56	7.71
	0.015	3.21	2.79	8.91	39.15	10.25	8.20
	0.030	3.24	3.60	11.84	40.10	9.89	11.08
30	0.000	3.16	4.70	10.71	43.01	11.15	7.20
	0.015	3.18	4.16	8.84	42.27	11.23	8.26
	0.030	3.20	4.64	9.73	43.23	10.75	9.53

Note: PI cost in $\times 10^{11}$ ($); gap values in %. All values are averages across five Beta distributions.

scenarios, achieving an average gap of 7.59% compared to 9.08% for DET. However, TPA consistently outperforms all benchmarks with an average gap of only 3.79%, demonstrating the robustness and consistency of the proposed approach across diverse demand patterns.

Under symmetric Beta distributions, the total costs produced by each model are relatively similar across parameter settings. This is because the cumulative realized demand tends to align closely with the cumulative nominal forecast over a long planning horizon T. In particular, the STATIC model yields an identical cost for each problem size I regardless of the distribution, because its decisions are entirely determined in Phase 1 before any demand realization. By contrast, more noticeable differences emerge under asymmetric Beta distributions. Under the Beta(0.5, 2) distribution, realized demand tends to fall below the nominal forecast, reducing the likelihood of stockouts and resulting in lower total costs across all models. Conversely, the Beta(2, 0.5) distribution generates demand realizations concentrated near the upper bound of the uncertainty set, increasing the risk of stockouts and leading to higher total costs. Notably, the TPA model maintains a relatively small optimality gap even under the high-demand Beta(2, 0.5) scenario, where the realized demand approaches the worst-case scenario that the model is designed to hedge against. This advantage extends to the comparison with SAA. Unlike TPA, SAA requires a specific distributional assumption and must be re-solved for each distribution. Despite this additional computational burden, TPA outperforms SAA across all 15 test instances, confirming that the proposed ARO approach provides hedging against distributional ambiguity.

To further assess the robustness of these findings under BOG inventory losses, we conduct sensitivity experiments with different BOG loss rates $\beta_j \in \{0, 0.015, 0.03\}$ for all three problem sizes ($I \in \{10, 20, 30\}$), where $\beta_j = 0.015$ and $\beta_j = 0.03$ correspond to 95% and 90% BOG scenarios, respectively. Table 5 reports the average PI cost and optimality gap across all five Beta distributions.

As shown in Table 5, the TPA framework exhibits the strongest robustness to BOG losses among all models across all problem sizes. The TPA optimality gap remains below 5% in all nine (I, β_j) combinations, and within each problem size, the gap variation is at most 1.53 percentage points ($I = 20$: 2.79–4.32%). In contrast, the DET model is most sensitive to BOG: its gap rises from 7.94% to 12.60% for $I = 10$ and from 8.78% to 11.84% for $I = 20$ as the loss rate increases from 0 to 0.03. The SAA benchmark also degrades under BOG losses (e.g., 7.20% → 10.30% for $I = 10$), whereas the SPOT and STATIC models remain insensitive due to their inherent inflexibility. Finally, the PI cost increases by less than 1.5% across all problem sizes when β_j rises from 0 to 0.03, confirming that the $\beta_j = 0$ baseline is a valid simplification for strategic planning.

Hereafter, we exclude the STATIC and SAA models from the remaining analyses because they differ fundamentally in their decision-making structure from the other approaches. STATIC commits all decisions before the realization of uncertainty (i.e., a here-and-now policy) and SAA requires re-optimization under a specific distributional assumption for each scenario. In contrast, the other benchmark models operate without such requirements. Figs. 2 and 3 present boxplots of total costs for 100 demand scenarios under symmetric and asymmetric distributions, respectively. Note that the SPOT model is omitted from these figures to preserve visual clarity, as its substantially higher costs obscure the comparison among the other models. The TPA model exhibits smaller variance than the DET model across all distributions, indicating robust performance under various types of demand uncertainty. No extreme outliers are observed, suggesting that none of the models produce cost failures under any scenario.

To summarize, the TPA framework demonstrates consistent near-optimal performance across all problem sizes and demand distributions examined in this section. Among the 15 test instances in Table 4, 14 instances achieve optimality gaps below 5%. The gaps range from 2.05–3.23% for $I = 10$ and from 4.07–4.99% for $I = 30$, with an overall average of 3.79%. The single exception occurs under the left-skewed Beta(0.5, 2) distribution at $I = 20$, where the gap reaches 6.34% as realized demand falls below the nominal forecast. Even in this case, TPA outperforms both DET (7.66%) and SAA (7.23%). The BOG sensitivity analysis (Table 5) further confirms that the TPA gap remains below 5% in all nine (I, β_j) combinations, with gap variation of at most 1.53% within each problem size, demonstrating robustness to physical inventory loss assumptions. These results confirm that the proposed TPA framework provides reliable near-optimal solutions across diverse problem scales, demand distributions, and BOG loss rates.

5.2. Cost structure and contract selection analysis

We now conduct a detailed analysis of the experimental results, focusing on individual cost components. To ensure computational

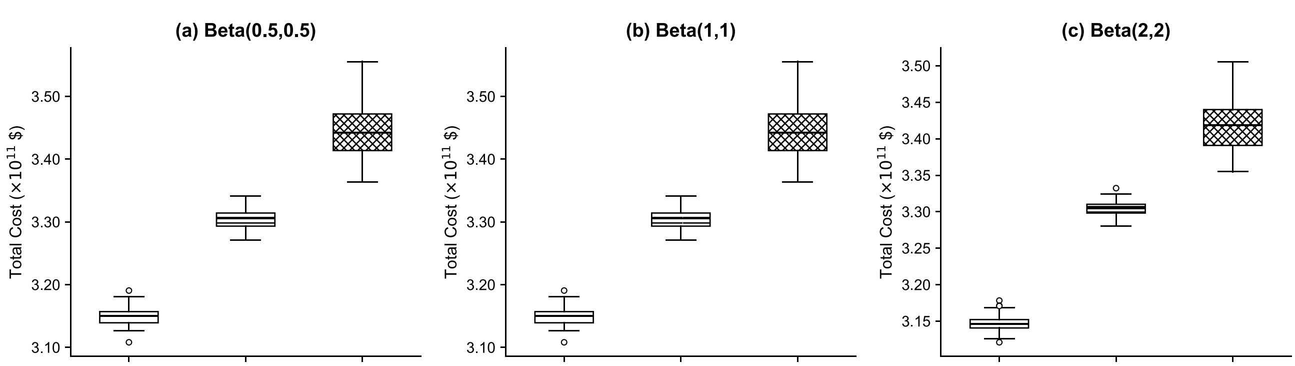

Fig. 2. Box plots for 100 total cost samples obtained from PI, TPA and DET model under symmetric distributions.

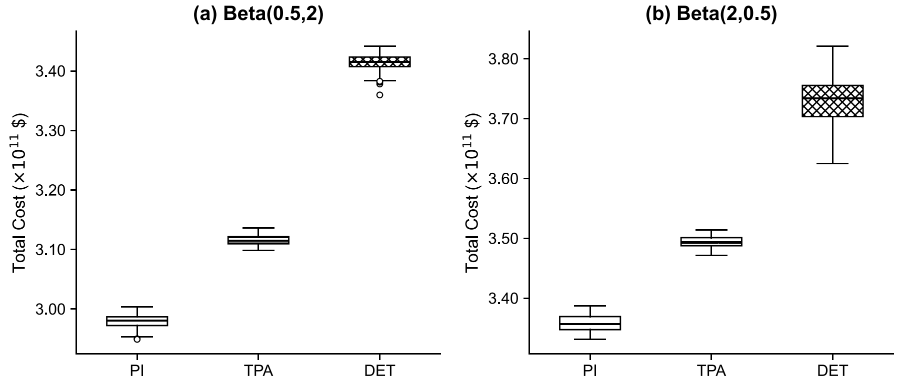

Fig. 3. Box plots for 100 total cost samples obtained from PI, TPA and DET models under asymmetric distributions.

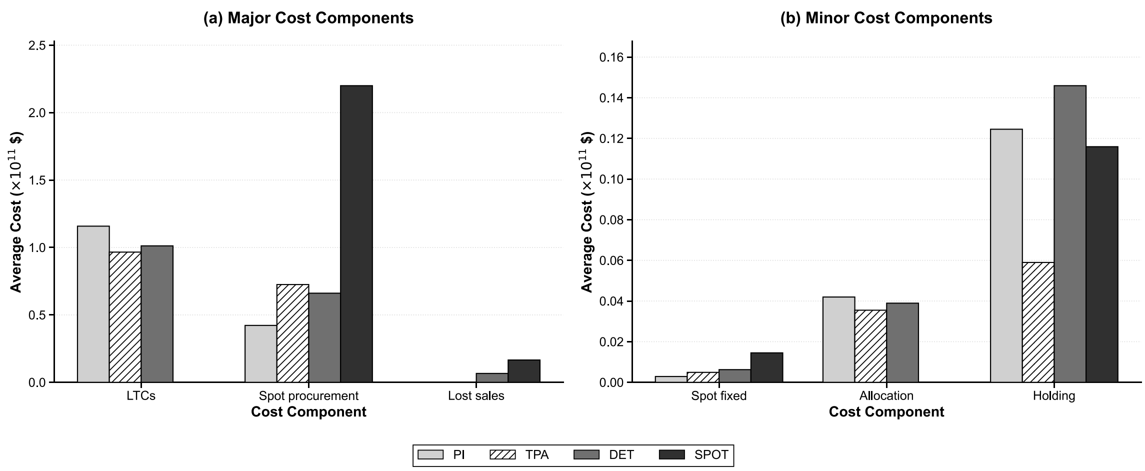

Fig. 4. Breakdown of individual cost components across models.

tractability for the extensive analyses that follow, we fix the number of contract candidates at $I = 10$ and the planning horizon at $T = 30$. In this section, we examine simulation results under the Beta(1, 1) distribution to understand how the TPA model constructs a more effective plan. Fig. 4 illustrates the breakdown of individual cost components for each model.

For the main cost components (Fig. 4(a)), the PI model allocates the largest portion of its budget to LTCs, followed by the DET model, with the TPA model allocating the least. This pattern arises because the PI model, with perfect foresight, can precisely match supply commitments to future demand realizations and thus confidently secures more volume through LTCs at lower unit prices. In contrast, the TPA model hedges against demand uncertainty by reserving flexibility for spot market procurement, resulting in lower LTC commitments. The SPOT model cannot utilize LTCs and must rely entirely on spot market procurement and lost sales to meet demand, resulting in the highest spot procurement cost. Comparing the TPA and DET models, the TPA model incurs slightly higher spot procurement costs than the DET model. However, the DET model experiences stockouts that lead to lost sales costs, which do not occur under the TPA model. This demonstrates that the TPA model effectively prevents demand shortfalls by strategically utilizing the spot market, whereas the DET model fails to meet demand in some scenarios.

Fig. 4(b) presents the minor cost components, which have a smaller cost scale than the major components. Among these, holding costs exhibit the most notable difference across models. The TPA model incurs substantially lower holding costs than the PI and DET models. This difference arises from the supply commitment structure inherent in LTCs: once a contract is signed, the contracted volume must be fully received in each period, as enforced by Constraint (15). Consequently, models that commit to larger LTC volumes (PI and DET) must store excess supply as inventory when demand falls below the contracted amount. The TPA model, by selecting a more moderate contract portfolio, avoids accumulating unnecessary inventory and achieves lower holding costs. This result demonstrates that contract selection derived from the DET model can lead to excessive inventory holdings, whereas the TPA model balances LTC commitments with flexible spot utilization.

We then examine the LTCs selected by each model. Fig. 5 compares the contract portfolios chosen by the TPA and DET models. The TPA and DET models consistently choose the same set of LTCs across all experiments. In contrast, the PI model occasionally selects different contracts depending on the specific demand scenario, as it optimizes decisions based on perfect future information. Therefore, we focus our comparison on the contract selections made by the TPA and DET models. Both models select contracts #3, #5, #8, and #9. However, a key difference emerges: the DET model selects contract #6 (active from periods 12 to 28), whereas the TPA model selects contract #7 (active from periods 15 to 30). Contract #6 delivers a larger quantity per period (1.98 million tons versus 1.59 million tons), while contract #7 provides a more moderate supply volume. This choice is consistent with the TPA model's strategy of selecting a moderate contract portfolio to reduce excess inventory, as discussed in Fig. 4(b).

Fig. 6 illustrates these inventory dynamics for periods 13–24 under a randomly generated demand scenario. The realized demand exhibits seasonality, with low demand in periods 13–16 and 21–22, and high demand in periods 17–20 and 23–24. Throughout both cycles, the DET model maintains elevated inventory levels, whereas the TPA model

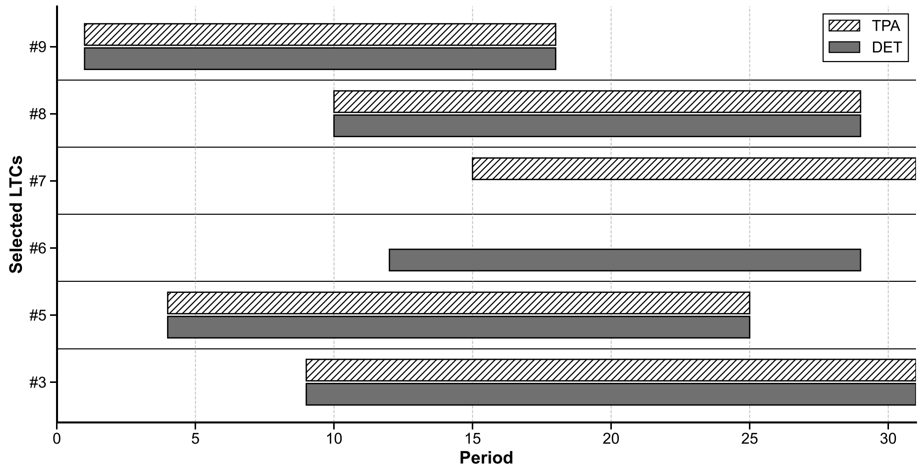

Fig. 5. Selected LTCs by the TPA and DET models.

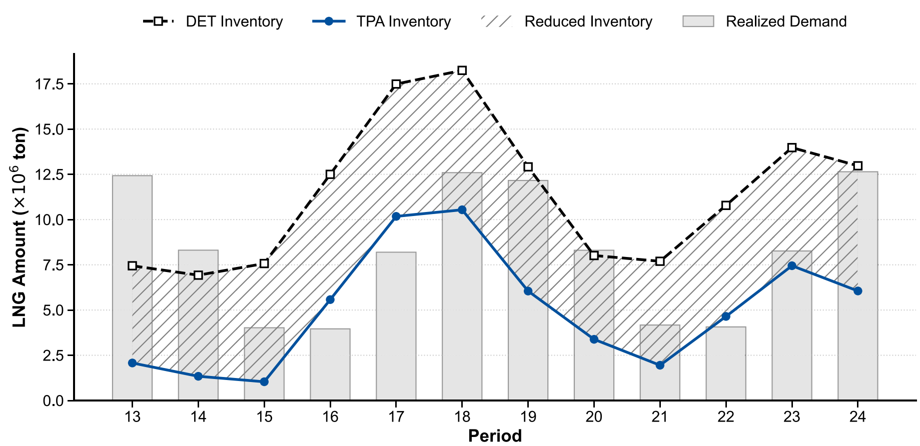

Fig. 6. Operational details for periods 13 to 24.

responds more adaptively to realized demand and consistently holds lower inventory. The hatched area represents the inventory reduction achieved by the TPA model relative to the DET model. This demonstrates that the TPA model dynamically maintains near-optimal inventory levels rather than adopting a fixed buffer irrespective of actual conditions.

5.3. Effects of forecast accuracy: Managing the size of the uncertainty set

This section investigates how forecast accuracy, controlled by the uncertainty parameter θ^{jt}, affects model performance. The parameter θ^{jt} determines the size of the uncertainty set: a smaller θ^{jt} represents more accurate forecasts with tighter demand bounds, while a larger θ^{jt} captures greater demand uncertainty. While our framework supports terminal- and time-specific uncertainty parameters θ^{jt} (Section 4), we set $\theta^{jt} = \theta$ for all $j \in \mathcal{J}$ and $t \in \mathcal{T}$ in these experiments to isolate the effect of overall forecast accuracy on model performance. In the baseline experiment, θ was set to 0.1, indicating demand realizations within $\pm 10\%$ of the nominal demand.

To comprehensively evaluate robustness across different forecasting environments, we conduct experiments with θ ranging from 0.05 to 0.50 in increments of 0.05. For each θ value, we re-optimize the TPA model, as this parameter directly affects the uncertainty set structure. We generate demand scenarios using five Beta distributions that capture diverse uncertainty patterns: bimodal (Beta(0.5, 0.5)), left-skewed (Beta(0.5, 2)), uniform (Beta(1, 1)), right-skewed (Beta(2, 0.5)), and bell-shaped (Beta(2, 2)). For each configuration, 100 demand scenarios are generated following the same procedure as in previous experiments. Table S2 of the Supplementary Material summarizes the resulting total costs and optimality gaps across all models and configurations, and Fig. 7 visualizes the optimality gaps for the TPA and DET models.

The results reveal a consistent pattern: the TPA model maintains substantially lower optimality gaps than the DET model across all uncertainty levels and distribution types. This performance advantage becomes more pronounced as θ increases, demonstrating the TPA model's robustness under higher demand uncertainty. Notably, the TPA model achieves optimality gaps within 5% for $\theta \leq 0.20$ across all distribution shapes. This threshold is practically relevant because recent advances in machine learning have significantly improved demand forecasting accuracy, with studies demonstrating statistically significant performance improvements over traditional methods [53]. Under such realistic forecasting conditions, the TPA model consistently delivers near-optimal performance regardless of the underlying demand distribution.

We now examine model performance across different distribution types. For symmetric distributions (uniform and bell-shaped), the DET model's gap increases dramatically from approximately 7% at $\theta = 0.05$ to over 40% at $\theta = 0.50$. In contrast, the TPA model's gap grows more moderately, from about 1% to 14%, representing a significant relative improvement. For asymmetric distributions, the performance differences are even more pronounced. Under the right-skewed distribution (Beta(2, 0.5)), where mean realized demand exceeds the forecasted demand, both total costs and optimality gaps escalate rapidly. At $\theta = 0.50$, the DET model exhibits a gap of 93%, indicating that its cost nearly doubles compared to the PI benchmark, whereas the TPA model maintains

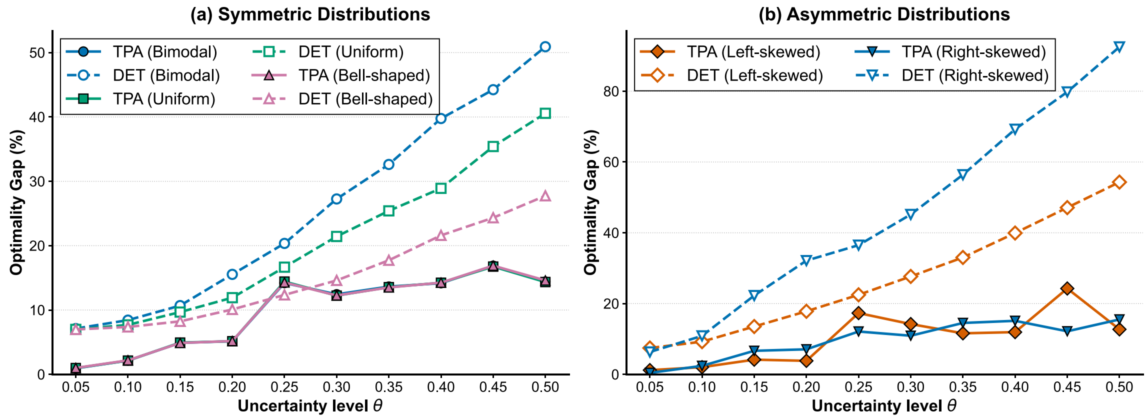

Fig. 7. Comparison of optimality gaps between TPA and DET across different uncertainty levels (θ) and demand distributions.

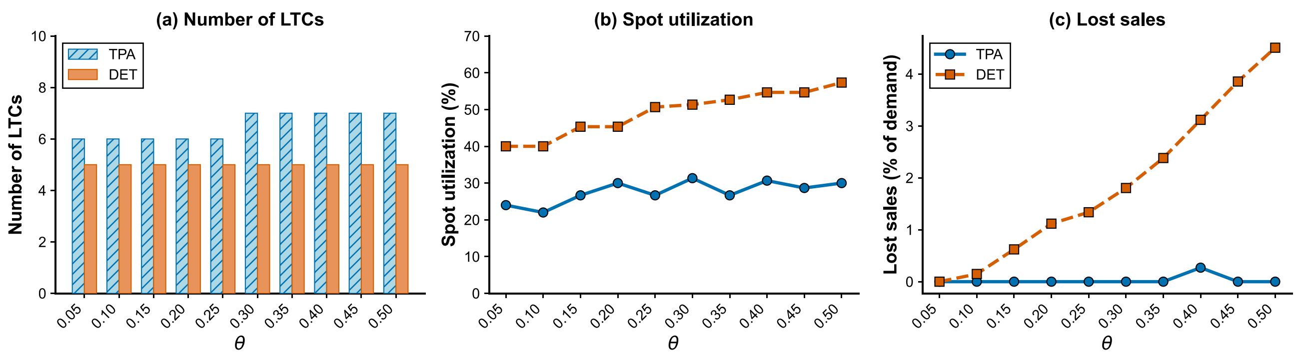

Fig. 8. Effect of the uncertainty parameter θ on the procurement portfolio: (a) number of LTCs selected, (b) spot market utilization, and (c) lost sales ratio.

a gap of only 16%, demonstrating its ability to adapt to upward demand deviations through flexible spot market utilization. In the case of the left-skewed distribution (Beta(0.5, 2)), the DET model exhibits a notably improved optimality gap compared to the right-skewed scenario. As indicated in Table 4, this is primarily because demand tends to be realized near the lower bound, enabling the system to satisfy requirements using existing inventory without incurring stockouts. However, even under these conditions, the TPA model consistently demonstrates superior performance compared to the DET model.

These results yield two key insights. First, the TPA model's adaptive framework provides effective protection against forecast errors across diverse demand patterns. While the DET model's performance degrades rapidly as uncertainty grows, the TPA model maintains near-optimal costs by dynamically adjusting operational decisions to realized demand. Second, the optimality gap between TPA and DET widens as forecast accuracy deteriorates, underscoring that RO becomes increasingly valuable in environments with high demand uncertainty. Conversely, when forecasts are highly accurate ($\theta = 0.05$), the DET approach can be a promising solution, as both models achieve similar performance levels.

Beyond cost performance, we examine how the uncertainty parameter θ affects the *portfolio structure* of procurement decisions. Fig. 8(a) shows that TPA adaptively increases the number of LTCs from 6 to 7 as θ grows from 0.05 to 0.50, securing additional contracted volume to hedge against a broader range of demand realizations. In contrast, DET maintains a fixed portfolio of 5 contracts regardless of θ, since its decisions are based solely on nominal demand forecasts without uncertainty awareness. Fig. 8(b) reveals the complementary adjustment: TPA's spot utilization fluctuates between 22% and 31%, reflecting its adaptive rebalancing between contracted and spot procurement as uncertainty widens. DET's spot utilization increases more steeply (40–57%), as it relies more heavily on spot purchases to offset the shortfall of its inflexible contract portfolio. This portfolio adjustment in response to the uncertainty level is the key mechanism behind TPA's cost advantage. By securing more LTCs under high uncertainty, TPA reduces its reliance on expensive spot procurement and mitigates lost sales penalties. Fig. 8(c) confirms this effect. DET's lost sales ratio increases steadily from near zero at $\theta = 0.05$ to 4.5% at $\theta = 0.50$, whereas TPA maintains negligible lost sales across the entire range (detailed results in Table S3 of the Supplementary Material).

5.4. Sensitivity analysis

We now examine how model performance varies with three key parameters: the periodic supply quantity (SUP), the spot market price (p_t), and the operational costs (a_{jt} and h_t). We introduce multipliers λ_{SUP}, λ_p, and λ_{ah} that scale the respective base values (i.e., values produced by Table 2). Each multiplier ranges from 0.2 to 2.0 in increments of 0.2, while other parameters remain at their base values. Fig. 9 reports the total cost of all four models across these configurations.

Fig. 9(a) shows the effect of λ_{SUP}. As the multiplier increases, the total costs of PI, TPA, and DET tend to decrease, whereas SPOT remains nearly constant. Higher SUP values increase the monthly supply volume U_{it} under LTCs, enhancing their cost-effectiveness relative to spot procurement. Because SPOT relies exclusively on spot procurement, it cannot exploit this benefit. Fig. 9(b) presents the sensitivity to λ_{ah}, which jointly scales allocation and holding costs. Total costs increase across all models as these operational costs rise. The DET model exhibits the steepest escalation compared to TPA and PI, consistent with our earlier finding that DET maintains higher inventory levels due to its deterministic contract selection. Fig. 9(c) depicts the impact of λ_p. SPOT displays a steep cost increase as the spot price grows, reflecting its complete dependence on the spot market and lost sales. The PI, TPA, and DET models exhibit a distinctive pattern: costs rise steeply for $\lambda_p \leq 0.6$ but grow more gradually beyond this threshold.

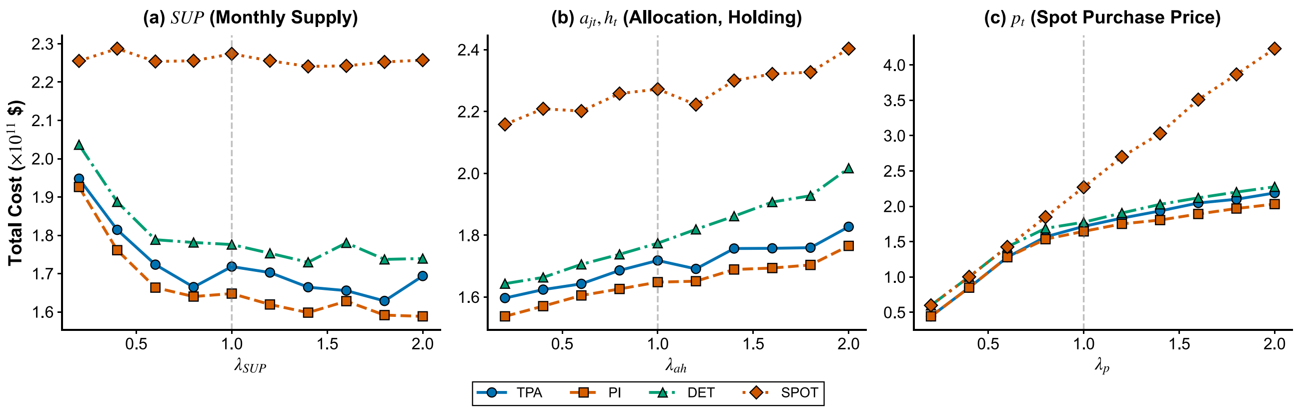

Fig. 9. Sensitivity of total cost to key parameter multipliers.

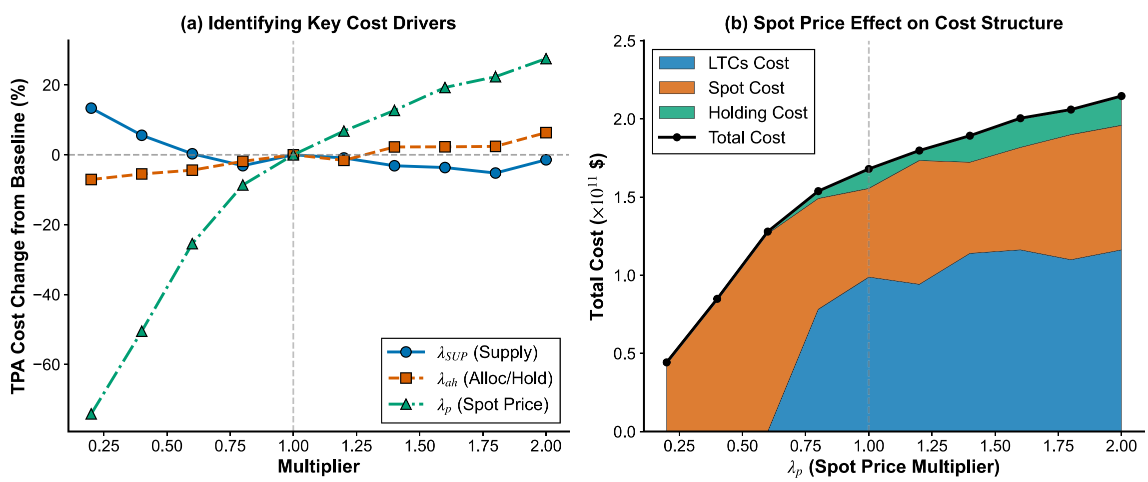

Fig. 10. Detailed cost breakdown under varying spot price multiplier λ_p.

In order to provide a detailed breakdown to explain this pattern, we present Fig. 10. Fig. 10(a) compares the percentage change in TPA total cost relative to the baseline setting ($\lambda = 1.0$) for each parameter. Among the three multipliers, λ_p exhibits the largest impact, ranging from approximately −65% at $\lambda_p = 0.2$ to +25% at $\lambda_p = 2.0$. In contrast, λ_{SUP} and λ_{ah} produce relatively modest variations of at most 15%. This result confirms that the spot price is the dominant cost driver for TPA.

Fig. 10(b) reveals how the cost structure shifts as λ_p increases. When $\lambda_p \leq 0.6$, spot procurement is substantially cheaper than LTCs, so TPA satisfies nearly all demand through the spot market. Beyond this threshold, LTCs become economically attractive, and their share of the total cost rises significantly. The holding cost also grows as TPA secures more contracted volume and carries additional inventory. These findings explain the pattern observed in Fig. 9(c). For $\lambda_p \leq 0.6$, total cost increases steeply because the model relies heavily on spot procurement whose price is rising. Beyond $\lambda_p = 0.6$, TPA increasingly utilizes LTCs, which mitigates the total cost increase by replacing spot market procurement with more stable procurement.

To further investigate how TPA adjusts its procurement strategy, we examine two metrics: the number of contracts $\left(\sum_{i\in\mathcal{I}} x_i\right)$ and spot utilization $\left(\sum_{j\in\mathcal{J},t\in\mathcal{T}} \delta_{jt}/(|\mathcal{J}|\cdot|\mathcal{T}|)\right)$. Fig. 11 illustrates these metrics under varying parameter multipliers.

Fig. 11(a) shows that as λ_{SUP} increases, both the number of contracts and spot utilization decrease. Higher monthly supply per contract enables TPA to satisfy demand with fewer LTCs, reducing dependence on the spot market. Fig. 11(b) reveals that beyond $\lambda_{ah} = 1.2$, the number of contracts decreases while spot utilization rises. LTCs tend to increase inventory and allocation volumes; when unit holding and allocation costs grow, TPA reduces contract usage and relies more on spot procurement to avoid these additional costs. Fig. 11(c) demonstrates that TPA begins utilizing LTCs at $\lambda_p \geq 0.8$, consistent with Fig. 10(b). Below this threshold, spot procurement remains cheaper, resulting in 100% spot utilization with no LTCs. As λ_p increases, TPA progressively adopts more LTCs to hedge against rising spot prices.

The preceding analysis examines sensitivity to cost parameters. We now investigate how *structural* market conditions—specifically, spot market accessibility—affect LNG supply chain planning. We introduce a spot capacity factor $\lambda_M \in [0.2, 1.5]$ that scales the baseline spot market capacity M, so that LNG procurement from the spot market is limited to $M \times \lambda_M$ per period (i.e., $w_{jt} \leq M \times \lambda_M$). Lower values of λ_M represent constrained market environments where spot procurement availability is restricted.

Fig. 12(a) reveals a significant difference in robustness. As λ_M decreases below 0.6, DET's optimality gap escalates dramatically—from 8% at $\lambda_M = 1.0$ to 585% at $\lambda_M = 0.2$—while TPA's gap remains stable at approximately 9% across all capacity levels. The right axis shows DET's lost sales reaching 45% of total demand at $\lambda_M = 0.2$, indicating severe supply disruption under constrained spot availability. TPA avoids this breakdown by adaptively adjusting its procurement strategy. As shown in Fig. 12(b), TPA's spot utilization ranges from 21% to 39%, whereas DET relies more heavily on the spot market (35–45%), reflecting its inability to hedge through contract flexibility. TPA's balanced procurement strategy prevents the persistent shortfalls that arise in DET (detailed results in Table S4 of the Supplementary Material).

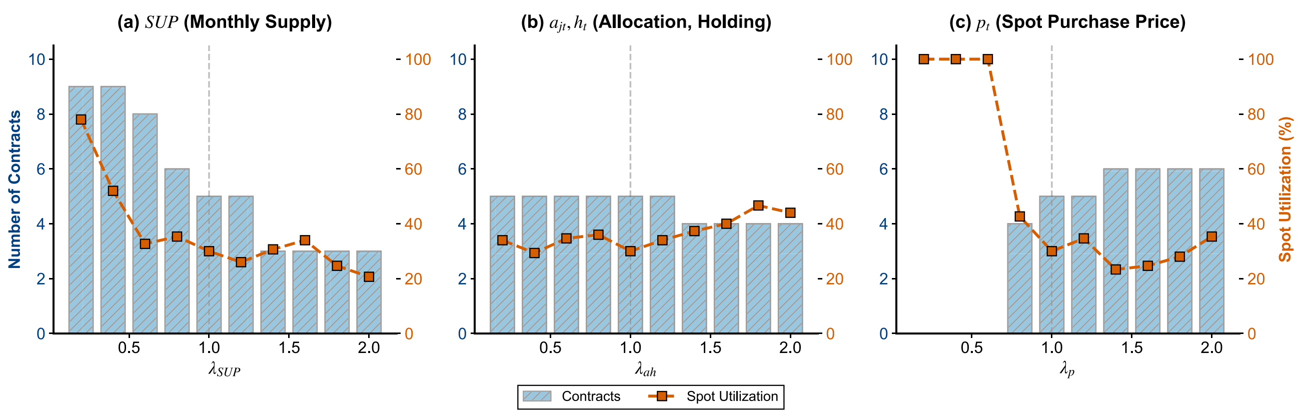

Fig. 11. Number of LTCs and spot utilization under varying parameter multipliers.

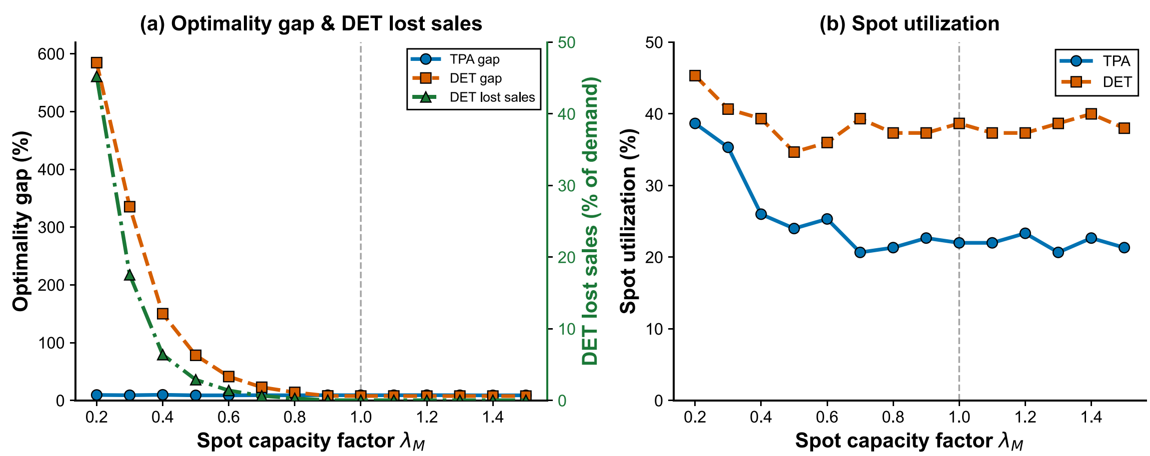

Fig. 12. Effect of spot market capacity on LNG supply chain planning: (a) optimality gap and DET lost sales ratio, and (b) spot utilization, under varying spot capacity factor λ_M.

To examine the interaction between spot capacity and spot price, we conduct a joint sensitivity analysis over a 6×7 grid of (λ_p, λ_M) combinations, where $\lambda_p \in \{0.6, 0.8, 1.0, 1.2, 1.5, 2.0\}$ and $\lambda_M \in \{0.2, 0.3, 0.5, 0.7, 1.0, 1.2, 1.5\}$. We define the *value of adaptivity* (VoA) as:

$$\text{VoA} = \frac{\text{Cost of DET} - \text{Cost of TPA}}{\text{Cost of TPA}} \times 100\%.$$

Fig. 13 presents the results as a two-dimensional sensitivity map. Fig. 13(a) shows that VoA reaches 797% under the most adverse conditions ($\lambda_M = 0.2$, $\lambda_p = 0.6$). VoA is driven primarily by capacity constraints rather than price: the bottom two rows ($\lambda_M \leq 0.3$) consistently exceed 200%, while the rows with $\lambda_M \geq 1.0$ remain below 12%. Figs. 13(b) and (c) reveal the underlying mechanism. DET's lost sales exceed 44% of total demand at $\lambda_M = 0.2$, indicating severe supply disruption, whereas TPA maintains lost sales below 3.8% across all conditions. This difference demonstrates that TPA expands the feasible operating range under adverse spot market conditions, providing robust performance even where DET's planning decisions lead to persistent shortfalls.

5.5. Managerial insights

The computational experiments confirm the practicality and strong performance of the TPA model in a real-world setting. Based on the cost component analysis, forecast accuracy experiments, and sensitivity analysis, we derive the following managerial insights for nationwide LNG supply chain planning.

- **Adaptive procurement and inventory management.** The TPA model secures stable supply commitments through LTCs while reserving spot market flexibility to absorb demand uncertainty. By incorporating demand observations into its LDRs, TPA dynamically adjusts inventory levels in response to realized demand, maintaining lower inventory than the DET model while ensuring supply reliability (Fig. 6). In contrast, the DET model selects contracts using only nominal demand forecasts, which can lead to stockouts and lost sales penalties when actual demand deviates from expectations.
- **Forecast accuracy and data-driven uncertainty calibration.** The TPA model achieves optimality gaps within 5% when $\theta \leq 0.20$; as forecast accuracy deteriorates, the performance gap between TPA and DET widens substantially, underscoring that RO becomes increasingly valuable under high demand uncertainty (Fig. 7). Our framework supports terminal- and time-specific uncertainty parameters θ^{jt}, accommodating the spatial and seasonal heterogeneity inherent in multi-terminal LNG systems. When historical forecast error data are available, we recommend setting $\theta^{jt} \approx$ 2–3 $\times$ CV_{jt}, where CV_{jt} denotes the CV of past forecast errors at terminal j in period t, providing a principled bound that covers the majority of historical forecast errors. Extreme fluctuations beyond this range (e.g., supply interruptions or geopolitical crises) are more effectively managed through dedicated disruption management approaches [54,55].
- **Spot price dominance and adaptive procurement response.** Among the parameters examined, spot price exhibits the largest impact on total cost, with variations ranging from −65% to +25% relative to the baseline (Fig. 10(a)), confirming it as the dominant cost driver. The TPA model adjusts its procurement strategy in

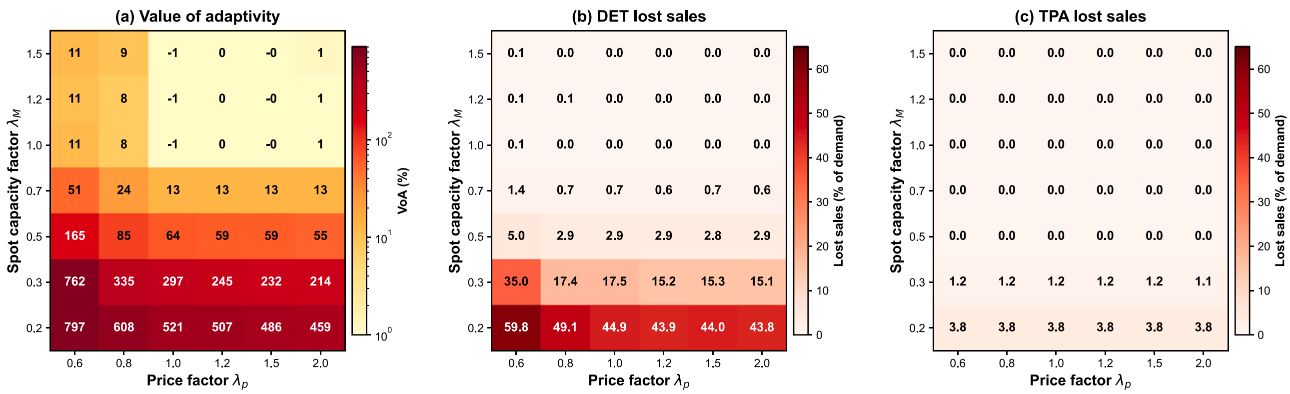

Fig. 13. Joint sensitivity analysis: (a) VoA (%), (b) DET lost sales (% of demand), and (c) TPA lost sales (% of demand) under varying spot capacity factor λ_M and price factor λ_p.

response to changing market conditions (Fig. 11). When spot prices rise above a threshold, TPA shifts toward LTCs to hedge against price volatility. Conversely, when allocation and holding costs increase, TPA reduces contract commitments and relies more on spot procurement to avoid inventory-related expenses.

- **Strategic resilience under spot market stress.** The structural sensitivity analysis reveals that TPA provides a robust hedge against spot market disruptions by maintaining feasible procurement plans even under severe capacity constraints. When spot capacity is severely constrained ($\lambda_M = 0.2$), TPA maintains a 9% optimality gap while DET's gap exceeds 585%, with lost sales reaching 45% of total demand (Fig. 12). The joint analysis shows that VoA increases nonlinearly as market conditions worsen, reaching 797% under the most adverse combination of constrained capacity and low spot prices (Fig. 13). This finding suggests that adaptive robust planning is most valuable precisely in tight spot market conditions, providing national energy planners with a quantitative justification for investing in optimization-based planning frameworks.
- **Energy security through penalty cost calibration.** National energy planners prioritizing supply security over cost minimization can achieve near-zero stockouts by setting the lost sales penalty l_{jt} sufficiently high relative to spot prices. When $l_{jt} \gg p_t$, the model allocates maximum supply through LTCs and spot procurement to eliminate demand shortfalls. This penalty-based approach allows governments to calibrate their acceptable level of supply disruption risk within the optimization framework.

6. Conclusions

This study addresses nationwide LNG supply chain planning from an optimization perspective. We consider four key decision areas: (i) LTC selection, (ii) spot market procurement, (iii) inter-terminal allocation, and (iv) inventory control. To efficiently capture both demand uncertainty and real-world applicability, we develop an ARO framework that effectively addresses these challenges. The framework explicitly considers forecast uncertainty and prescribes policies that adjust dynamically as actual demand is realized—capabilities that traditional, forecast-driven plans lack.

Our solution builds on a customized two-phase approach that transforms a stochastic optimization model into two tractable subproblems. Extensive computational experiments confirm the effectiveness and applicability of the proposed framework. The framework achieves an optimality gap within 5% relative to a perfect information lower bound, substantially outperforming the industry-standard deterministic benchmark. Moreover, the sensitivity analysis reveals that the spot price is the dominant cost driver, with cost variations ranging from −65% to +25% relative to the baseline. The framework demonstrates strategic adaptability by shifting between LTCs and spot procurement in response to environmental changes. Structural sensitivity analysis further confirms that TPA provides a robust hedge against spot market disruptions (under severely constrained conditions ($\lambda_M = 0.2$)). TPA maintains a 9% optimality gap while the deterministic benchmark deteriorates sharply (585% gap with lost sales reaching 45% of total demand). The VoA increases under adverse conditions, reaching 797%, demonstrating that adaptive robust planning is most valuable precisely in tight spot market conditions.

In summary, the proposed adaptive robust framework secures a larger share of supply through LTCs while utilizing the spot market and inventory more flexibly. The framework provides reliable and economical nationwide planning guidelines. To the best of our knowledge, this study is the first to formulate national LNG supply chain planning as an adaptive decision problem and to demonstrate its tractability on realistic data. Because the framework delivers clear, implementable policies, it can be a practical decision-support tool for long-range LNG supply chain planners.

We conclude by discussing the scope of applicability of the proposed framework and directions for future research. We note that the reported near-optimality results (within 5% of the perfect information bound) are predicated on three modeling conditions that define the framework's scope of applicability. First, demand uncertainty is characterized through interval-based uncertainty sets centered on forecasts, enabling exact reformulation via strong duality. Second, cost functions are modeled as linear per-unit rates, which are well-suited to strategic planning where costs are aggregated over bi-monthly periods; because both the cost structure and the decision rules are linear in decision variables, the robust counterpart can be reformulated as a tractable linear program. Third, the framework assumes normal market conditions without structural disruptions, which are more appropriately addressed through disruption management approaches (discussed below). When any of these conditions is significantly violated, the near-optimality guarantee may not directly transfer.

Several avenues for future research remain. First, while demand uncertainty is fully incorporated, the spot price is treated as a deterministic parameter due to a fundamental structural reason: if p_t were also uncertain, bilinear terms $p_t \times d^{j'k}$ would invalidate the strong duality reformulation [15,18]. Jointly modeling demand and price uncertainty requires fundamentally different approaches and remains a priority direction (see also Appendix B). Second, the current contract selection logic does not account for qualitative factors such as geopolitical risk, supplier reliability, and regional diversification. Structural disruptions—such as supply interruptions from geopolitical conflicts, extreme price spikes from market regime shifts, or sudden infrastructure failures—fall outside the scope of interval-based uncertainty sets. Pre-emptively hedging against such low-probability, high-impact events leads to overly conservative solutions. These scenarios are more effectively addressed through disruption management [54,55], which focuses on rapid recovery, alternative resource activation, and supply chain resilience. Third, incorporating cargo diversion and resale options would provide additional flexibility mechanisms for LNG supply chain planning. Additionally, when sufficient terminal-level historical demand data become available, distributionally robust optimization with moment-based or Wasserstein-based ambiguity sets could provide a richer characterization of demand uncertainty while retaining computational tractability. Fourth, extending the cost model to incorporate nonlinear structures, such as congestion effects or economies of scale, would broaden the framework's applicability to operational-level planning, requiring solution approaches beyond LDRs. Finally, a comprehensive comparison with decomposition-based multistage stochastic programming methods such as scenario-tree approaches remains a valuable direction for future research. We hope this study will enhance future research on nationwide LNG supply chain strategies and inspire ongoing investigations into robust, multistage energy logistics optimization.

CRediT authorship contribution statement

Euihyeon Choi: Writing – original draft, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Junhyeok Lee:** Writing – original draft, Visualization, Validation, Software, Methodology. **Ilkyeong Moon:** Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors are grateful for the valuable comments from the associate editor and anonymous reviewers. This work was supported by the National Research Foundation of Korea (NRF) grants funded by the Korean government (Ministry of Science and ICT) [Grant Nos. RS-2023-00218913 and RS-2024-00337285].

Appendix A. Mathematical notation

Table A.6
Parameters for the mathematical model.

Symbol	Description
Indices and sets	
$i \in \mathcal{I}$	Set of LTC candidates, $\mathcal{I} = \{1, \dots, I\}$
$j \in \mathcal{J}$	Set of LNG terminals, $\mathcal{J} = \{1, \dots, J\}$
$t \in \mathcal{T}$	Set of time periods, $\mathcal{T} = \{1, \dots, T\}$
$t \in \mathcal{T}^+$	Extended time period set for inventory, $\mathcal{T}^+ = \{1, \dots, T+1\}$
D^{jt}	Uncertainty set for demand at terminal j in period t
$\mathbf{D}^t$	Collection of uncertainty sets up to period t, $(D^{jk}, j \in \mathcal{J}, k \in \{1, \dots, t\})$
$\mathbf{D}(\gamma)$	Adjustable uncertainty set parameterized by γ, $\mathbf{D}^T(\gamma)$
Contract-related parameters	
SUP_i	Periodic supply amount for contract i (tonnes)
START_i	Contract start period for contract i
END_i	Contract end period for contract i
PRICE_i	Unit price per tonne for contract i
FC_i	Fixed cost for contract i
C_i	Total cost of contract i, $\mathrm{FC}_i + \mathrm{SUP}_i \times \mathrm{PRICE}_i \times (\mathrm{END}_i - \mathrm{START}_i + 1)$
U_{it}	Supply available from contract i in period t (SUP_i if active, 0 otherwise)
Cost parameters	
a_{jt}	Per-tonne allocation cost for terminal j in period t
p_t	Spot market unit price in period t
f_{jt}	Fixed cost for activating spot market at terminal j in period t
h_t	Per-tonne inventory holding cost in period t
l_{jt}	Per-tonne lost sales penalty cost at terminal j in period t
Capacity and BOG parameters	
$\bar{s}_j$	Maximum storage capacity at terminal j (tonnes)
$\underline{s}_j$	Minimum inventory level (safety stock) at terminal j (tonnes)
M	Maximum possible quantity of LNG procured from the spot market (tonnes)
β_j	Periodic BOG loss rate at terminal j, $\beta_j \in [0, 1)$
Demand-related parameters	
$\tilde{d}^{jt}$	Random demand at terminal j in period t
d^{jt}	Realized demand at terminal j in period t
$\hat{d}^{jt}$	Nominal/forecasted demand at terminal j in period t
$\underline{d}^{jt}$	Lower bound of demand at terminal j in period t, $\hat{d}^{jt} - b^{jt}$
$\overline{d}^{jt}$	Upper bound of demand at terminal j in period t, $\hat{d}^{jt} + b^{jt}$
b^{jt}	Demand deviation bound, $\theta^{jt}\hat{d}^{jt}$
θ^{jt}	Demand uncertainty parameter at terminal j in period t, $\theta^{jt} \in [0, 1]$
$\tilde{\mathbf{d}}^t$	Random vector of demands from periods 1 through t
$\tilde{\mathbf{d}}$	Random vector of all demands over planning horizon, $\tilde{\mathbf{d}}^T$
$\mathbf{d}^t$	Realized demand vector through period t
$\mathbf{d}$	Realized demand vector for entire planning horizon
$\check{\mathbf{d}}(\gamma)$	Worst-case scenario of uncertainty (WSU), $(\hat{d}^{jt} + \gamma b^{jt}, \forall j \in \mathcal{J}, t \in \mathcal{T})$
Control parameters	
ξ	Cost target in TRO (Phase 1)
γ	Variable for determining the uncertainty set size, $\gamma \in [0, 1]$
α	Target coefficient, $\alpha \in [0, 1]$
$\rho(\gamma)$	Optimal objective value of $(\mathrm{P}_{\rho(\gamma)})$ parameterized by γ

Table A.7
Decision variables of the mathematical model.

Symbol	Description
Binary decision variables	
x_i	1 if LTC i is signed, 0 otherwise
δ_{jt}	1 if spot market is activated at terminal j in period t, 0 otherwise
Continuous decision variables	
y_{jt}	Amount of LNG allocated to terminal j in period t (tonnes)
s_{jt}	Initial inventory at terminal j at start of period t (tonnes)
r_{jt}	Amount of inventory used to fulfill demand at terminal j in period t (tonnes)
w_{jt}	Quantity of LNG procured from spot market at terminal j in period t (tonnes)
z_{jt}	Stockout (lost sales) quantity of LNG at terminal j in period t (tonnes)
Binary adjustable variables	
$\delta_{jt}(\tilde{\mathbf{d}}^{t-1})$	1 if spot market is activated at terminal j at the start of period t, 0 otherwise
$\boldsymbol{\delta}(\tilde{\mathbf{d}})$	Collection of adjustable binary decisions, $(\delta_{jt}(\tilde{\mathbf{d}}^{t-1}), \forall j \in \mathcal{J}, t \in \mathcal{T})$
Continuous adjustable variables (information up to $t-1$)	
$y_{jt}(\tilde{\mathbf{d}}^{t-1})$	Allocation decision at terminal j at the start of period t
$s_{jt}(\tilde{\mathbf{d}}^{t-1})$	Inventory level at terminal j at the start of period t
$\boldsymbol{\mu}(\tilde{\mathbf{d}})$	Collection of adaptive decisions based on demand up to $t-1$: $(y_{jt}(\tilde{\mathbf{d}}^{t-1}), s_{jt}(\tilde{\mathbf{d}}^{t-1}))$
Continuous adjustable variables (information up to t)	
$r_{jt}(\tilde{\mathbf{d}}^{t})$	Fulfillment quantity at terminal j at the end of period t
$w_{jt}(\tilde{\mathbf{d}}^{t})$	Spot procurement quantity at terminal j at the end of period t
$z_{jt}(\tilde{\mathbf{d}}^{t})$	Stockout (lost sales) quantity at terminal j at the end of period t
$\boldsymbol{\nu}(\tilde{\mathbf{d}})$	Collection of adaptive decisions based on demand up to t: $(r_{jt}(\tilde{\mathbf{d}}^{t}), w_{jt}(\tilde{\mathbf{d}}^{t}), z_{jt}(\tilde{\mathbf{d}}^{t}))$
Combined adjustable variables	
$\boldsymbol{\pi}(\mathbf{d})$	Set of all adjustable variables, $(\boldsymbol{\delta}(\mathbf{d}), \boldsymbol{\mu}(\mathbf{d}), \boldsymbol{\nu}(\mathbf{d}))$
LDR coefficients	
y^0_{jt}	Constant coefficient for allocation LDR
$y^{j'k}_{jt}$	Demand-dependent coefficient for allocation LDR
s^0_{jt}	Constant coefficient for inventory LDR
$s^{j'k}_{jt}$	Demand-dependent coefficient for inventory LDR
r^0_{jt}	Constant coefficient for fulfillment LDR
$r^{j'k}_{jt}$	Demand-dependent coefficient for fulfillment LDR
w^0_{jt}	Constant coefficient for spot procurement LDR
$w^{j'k}_{jt}$	Demand-dependent coefficient for spot procurement LDR
z^0_{jt}	Constant coefficient for stockout (lost sales) LDR
$z^{j'k}_{jt}$	Demand-dependent coefficient for stockout (lost sales) LDR

Appendix B. Model extensions for operational constraints

This appendix presents extension formulations for operational constraints that may be relevant in different LNG supply chain contexts. While these constraints are non-binding for South Korea's strategic-level planning (due to sufficient infrastructure capacity relative to import volumes), they can be readily incorporated for applications to infrastructure-constrained systems or operational-level scheduling models.

Cargo diversion and resale. For markets with limited storage capacity or active secondary LNG trading, excess LTC supply can be diverted and resold rather than being fully allocated to domestic terminals. This extension ensures model feasibility when storage capacity cannot absorb all contracted volumes. To this end, we introduce $u_{it} \geq 0$ representing the quantity of LNG procured from LTC candidate i and diverted for resale in period t. The objective function is modified to include resale revenue: Objective function (1) $- \sum_{i\in\mathcal{I}} \sum_{t\in\mathcal{T}} p_t^{\text{resale}} u_{it}$ where p_t^{resale} is the resale price in period t. We note that many Korean LTCs contain destination clauses that restrict cargo diversion to alternative buyers or regions. Achieving resale flexibility under such contracts would require either renegotiating existing terms or procuring destination-free contracts, typically at a price premium.

LTC swing clauses. This modification requires introducing the variable y_{ijt} to track allocation from each contract to each terminal, with the original allocation variable becoming $y_{jt} = \sum_{i\in\mathcal{I}} y_{ijt}$. For markets where contracts include swing provisions (typically ±5–15% volume flexibility), the LTC allocation constraint (2) can be modified as:

$$(1-\eta_i)U_{it}x_i \leq \sum_{j\in\mathcal{J}} y_{ijt} + u_{it} \leq (1+\eta_i)U_{it}x_i, \qquad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{B.59}$$

where $\eta_i \in [0, 0.15]$ represents the swing flexibility percentage for contract i.

Take-or-pay obligations. The full LTC allocation constraint (2) in our base model ($\mathrm{P_{DET}}$) already captures take-or-pay obligations, requiring all contracted volumes to be allocated each period. For partial take-or-pay flexibility (e.g., 85% minimum), the constraint can be relaxed to:

$$\sum_{j\in\mathcal{J}} y_{ijt} + u_{it} \geq \kappa U_{it} x_i, \qquad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{B.60}$$

where $\kappa \in [0, 1]$ is the take-or-pay percentage.

Berth/throughput and regasification capacity. Terminal receiving and throughput capacity can be modeled by adding the following constraint:

$$\sum_{i\in\mathcal{I}} y_{ijt} \leq B_j, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{B.61}$$

where B_j represents the receiving and throughput capacity at terminal j per period. For systems where regasification may be constrained, the following constraint can be added:

$$r_{jt} \leq R_j, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{B.62}$$

where R_j is the regasification capacity at terminal j per period.

Spot procurement lead time. For applications with shorter planning periods (e.g., days or weeks) where LNG cargo lead times cannot be absorbed within a single period, explicit lead time modeling can be incorporated. The spot arrival quantity at terminal j in period t becomes:

$$g_{jt} = w_{j,t-\ell}, \qquad \forall j \in \mathcal{J}, t \in \mathcal{T} \tag{B.63}$$

where ℓ represents the procurement lead time in periods (typically 2–4 weeks for Northeast Asian destinations). The inventory balance equation is then modified to use g_{jt} instead of w_{jt}, ensuring that spot orders placed in period $t-\ell$ arrive in period t. In our bi-monthly framework, each period (approximately 60 days) is sufficiently long to accommodate typical LNG voyage times (17–34 days), so this extension is not required.

Demand-correlated spot pricing. In LNG markets with strong seasonal demand–price co-movement, spot prices tend to spike during high-demand periods (e.g., winter heating seasons in Northeast Asia). Rather than treating p_t as independent of demand, the spot price can be calibrated as a function of forecasted aggregate demand:

$$p_t = \varphi \sum_{j\in\mathcal{J}} \hat{d}^{jt},$$

where $\varphi > 0$ is a scaling coefficient calibrated from historical spot price–demand data. Under our bi-monthly seasonal demand structure (with a winter-to-summer ratio of approximately 3:1), this formulation automatically generates higher spot prices in high-demand periods, capturing the observed demand–price correlation. Importantly, p_t remains a deterministic parameter derived from the nominal demand forecast, so the robust optimization structure, including the tractable LP reformulation via strong duality, is fully preserved. This extension complements the independently sampled spot prices used in our computational experiments (Section 5), which are designed to test algorithmic robustness across diverse price–demand configurations.

Unified extended model. The extensions above can be combined into a unified extended deterministic model $\mathrm{P_{EXT}}$ that incorporates all operational constraints. This formulation disaggregates allocation decisions by contract (y_{ijt} replacing y_{jt}) to enable swing clause modeling, introduces the resale variable u_{it} for cargo diversion flexibility, and explicitly models spot procurement lead times through the arrival variable g_{jt}.

($\mathrm{P}_{\mathrm{EXT}}$)

$$\min \sum_{i\in\mathcal{I}} C_i x_i + \sum_{t\in\mathcal{T}} \left(\sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}} a_{jt} y_{ijt} + \sum_{j\in\mathcal{J}} f_{jt}\delta_{jt} + \sum_{j\in\mathcal{J}} p_t w_{jt} + \sum_{j\in\mathcal{J}} l_{jt} z_{jt} + \sum_{j\in\mathcal{J}} h_t s_{jt+1} \right) - \sum_{i\in\mathcal{I}} \sum_{t\in\mathcal{T}} p_t^{\mathrm{resale}} u_{it} \tag{B.64}$$

$$\text{s.t.} \quad (1-\eta_i)U_{it}x_i \le \sum_{j\in\mathcal{J}} y_{ijt} \le (1+\eta_i)U_{it}x_i, \quad \forall i \in \mathcal{I}, t\in\mathcal{T} \tag{B.65}$$

$$\sum_{j\in\mathcal{J}} y_{ijt} + u_{it} \ge \kappa U_{it} x_i, \quad \forall i \in \mathcal{I}, t\in\mathcal{T} \tag{B.66}$$

$$\sum_{i\in\mathcal{I}} y_{ijt} \le B_j, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.67}$$

$$r_{jt} \le R_j, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.68}$$

$$s_{jt} + \sum_{i\in\mathcal{I}} y_{ijt} \le \bar{s}_j, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.69}$$

$$\underline{s}_j \le s_{jt}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T}^+ \tag{B.70}$$

$$s_{jt+1} = (1-\beta_j)s_{jt} + \sum_{i\in\mathcal{I}} y_{ijt} - r_{jt}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.71}$$

$$g_{jt} = w_{j,t-\ell}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} : t > \ell \tag{B.72}$$

$$r_{jt} + g_{jt} + z_{jt} = d^{jt}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.73}$$

$$w_{jt} \le M\delta_{jt}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.74}$$

$$y_{ijt}, w_{jt}, r_{jt}, z_{jt}, g_{jt}, u_{it} \ge 0, \quad \forall i\in\mathcal{I}, j \in \mathcal{J}, t\in\mathcal{T} \tag{B.75}$$

$$x_i \in \{0,1\}, \quad \forall i \in \mathcal{I}; \quad \delta_{jt} \in \{0,1\}, \quad \forall j \in \mathcal{J}, t\in\mathcal{T} \tag{B.76}$$

where the additional parameters are: $\eta_i \in [0, 0.15]$ (swing flexibility percentage for contract i), $\kappa \in [0,1]$ (take-or-pay minimum percentage), B_j (berth/receiving capacity), R_j (regasification capacity), and ℓ (spot procurement lead time in periods). Constraint (B.72) links spot orders to arrivals with lead time ℓ (i.e., $g_{jt} = 0$ for $t \le \ell$). The base model $\mathrm{P}_{\mathrm{DET}}$ is recovered by setting $\eta_i = 0$, $\kappa = 1$, $u_{it} = 0$, $\ell = 0$, and $B_j, R_j \to \infty$ for all i, j, t.

Appendix C. Supplementary data

Supplementary data for this article can be found online at doi:10.1016/j.apenergy.2026.127670.

Data availability

Data will be made available on request.

References

[1] Shell International Ltd. Shell LNG Outlook 2025, Industry report, Shell International Ltd.; 2025. https://www.shell.com/what-we-do/oil-and-natural-gas/liquefied-natural-gas-lng/lng-outlook-2025 [accessed: 14 May 2025].

[2] GIIGNL. GIIGNL 2024 Annual Report, Annual report, International Group of Liquefied Natural Gas Importers; 2024. https://giignl.org/giignl-releases-2024-annual-report/ [accessed: 14 May 2025].

[3] Nikhalat-Jahromi H, Angeloudis P, Bell MGH, Cochrane RA. Global LNG trade: a comprehensive up to date analysis. Marit Econ Logist 2017;19:160–81.

[4] Hupka Y, Popova I, Simkins B, Lee T. A review of the literature on LNG: hubs development, market integration, and price discovery. J Commer Mark 2023;31:100349.

[5] Pradipta MA. Investigation of the price linkage between Asian LNG spot and far east Asian LNG prices and its implications. Indones J Energy 2018;1(1):52–65.

[6] Jayadev G, Leibowicz BD, Bard JF, Çalcı B. Strategic interactions between liquefied natural gas and domestic gas markets: a bilevel model. Comput Oper Res 2022;144:105807.

[7] Khalilpour R, Karimi IA. Selection of liquefied natural gas (LNG) contracts for minimizing procurement cost. Ind Eng Chem Res 2011;50(17):10298–312.

[8] Geng J-B, Ji Q, Fan Y, Shaikh F. Optimal LNG importation portfolio considering multiple risk factors. J Clean Prod 2017;151:452–64.

[9] Shahrukh M, Srinivasan R, Karimi IA. Optimal procurement of liquefied natural gas cargos from long-term contracts and spot market through mathematical programming. Ind Eng Chem Res 2021;60(9):3658–69.

[10] Biresselioglu ME, Demir MH, Kandemir C. Modeling turkey's future LNG supply security strategy. Energy Policy 2012;46:144–52.

[11] Sangaiah AK, Tirkolaee EB, Goli A, Dehnavi-Arani S. Robust optimization and mixed-integer linear programming model for LNG supply chain planning problem. Soft comput 2020;24:7885–905.

[12] Eriksen U, Kristiansen J, Fagerholt K, Pantuso G. Planning a maritime supply chain for liquefied natural gas under uncertainty. Marit Transp Res 2022;3:100061.

[13] Al-Haidous S, Govindan R, Elomri A, Al-Ansari T. An optimization approach to increasing sustainability and enhancing resilience against environmental constraints in LNG supply chains: a qatar case study. Energy Rep 2022;8:9742–56.

[14] Li M, Fagerholt K, Schütz P. Maritime inventory routing with transshipment: the case of yamal LNG. Flex Serv Manuf J 2023;35(1):269–94.

[15] Ben-Tal A, El Ghaoui L, Nemirovski A. Robust optimization, vol. 28. Princeton University Press; 2009.

[16] Soyster AL. Convex programming with set-inclusive constraints and applications to inexact linear programming. Oper Res 1973;21(5):1154–7.

[17] Ben-Tal A, Nemirovski A. Robust solutions of uncertain linear programs. Oper Res Lett 1999;25(1):1–13.

[18] Yanıkoğlu İ, Gorissen BL, Den Hertog D. A survey of adjustable robust optimization. Eur J Oper Res 2019;277(3):799–813.

[19] Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A. Adjustable robust solutions of uncertain linear programs. Math Program 2004;99(2):351–76.

[20] Bertsimas D, Thiele A. A robust optimization approach to inventory theory. Oper Res 2006;54(1):150–68.

[21] See C-T, Sim M. Robust approximation to multiperiod inventory management. Oper Res 2010;58(3):583–94.

[22] Bertsimas D, Iancu DA, Parrilo PA. A hierarchy of near-optimal policies for multistage adaptive optimization. IEEE Trans Autom Control 2011;56(12):2809–24.

[23] Georghiou A, Wiesemann W, Kuhn D. Generalized decision rule approximations for stochastic programming via liftings. Math Program 2015;152:301–38.

[24] Bertsimas D, Goyal V, Lu y B. A tight characterization of the performance of static solutions in two-stage adjustable robust linear optimization. Math Program 2015;150(2):281–319.

[25] Lim YF, Wang C. Inventory management based on target-oriented robust optimization. Manag Sci 2017;63(12):4409–27.

[26] Marandi A, Den Hertog D. When are static and adjustable robust optimization problems with constraint-wise uncertainty equivalent? Math Program 2018;170:555–68.

[27] Hanasusanto GA, Kuhn D, Wiesemann W. K-adaptability in two-stage robust binary programming. Oper Res 2015;63(4):877–91.

[28] Bertsimas D, Dunning I. Multistage robust mixed-integer optimization with adaptive partitions. Oper Res 2016;64(4):980–98.

[29] Postek K, Den Hertog D. Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. INFORMS J Comput 2016;28(3):553–74.

[30] Bertsimas D, Georghiou A. Binary decision rules for multistage adaptive mixed-integer optimization. Math Program 2018;167:395–433.

[31] Ben-Tal A, Golany B, Nemirovski A, Vial J-P. Retailer-supplier flexible commitments contracts: a robust optimization approach. Manuf Serv Oper Manag 2005;7(3):248–71.

[32] Kim BS, Chung do B. Affinely adjustable robust model for multiperiod production planning under uncertainty. IEEE Trans Eng Manag 2017;64(4):505–14.

[33] Lim YF, Jiu S, Ang M. Integrating anticipative replenishment allocation with reactive fulfillment for online retailing using robust optimization. Manuf Serv Oper Manag 2021;23(6):1616–33.

[34] Jiu S. Robust omnichannel retail operations with the implementation of ship-from-store. Transp Res Part E Logist Transp Rev 2022;157:102550.

[35] Lee J, Moon I. A decomposition approach for robust omnichannel retail operations considering the third-party platform channel. Transp Res Part E Logist Transp Rev 2024;184:103466.

[36] Riepin I, Schmidt M, Baringo L, Müsgens F. Adaptive robust optimization for European strategic gas infrastructure planning. Appl Energy 2022;324:119686.

[37] Gu C, Tang C, Xiang Y, Xie D. Power-to-gas management using robust optimisation in integrated energy systems. Appl Energy 2019;236:681–9.

[38] Shen F, Zhao L, Du W, Zhong W, Qian F. Large-scale industrial energy systems optimization under uncertainty: a data-driven robust optimization approach. Appl Energy 2020;259:114199.

[39] Abdin AF, Caunhye AM, Zio E, Cardin M-A. Optimizing generation expansion planning with operational uncertainty: a multistage adaptive robust approach. Appl Energy 2022;306:118032.

[40] Bazyar A, Zarrinpoor N, Safavian A. Optimal design of a sustainable natural gas supply chain network under uncertainty. Chem Eng Res Des 2021;176:60–88.

[41] Keyvandarian A, Saif A. An adaptive distributionally robust optimization approach for optimal sizing of hybrid renewable energy systems. J Optim Theory Appl 2024;203(2):2055–82.

[42] Fulwood M. LNG Shipping Chokepoints, Energy Insight NG 188, Oxford Institute for Energy Studies. Feb 2024. https://www.oxfordenergy.org/wpcms/wp-content/uploads/2024/02/NG-188-LNG-Shipping-Chokepoints.pdf.

[43] Zhang J, Yin X, Lei Z, Wang J, Fan Z, Liu S. Economic feasibility of LNG business: an integrated model and case study analysis. Energies 2024;17:3351.
[44] Institute for Energy Economics and Financial Analysis. South Korea's LNG Overbuild, Briefing note, IEEFA. Nov 2023. https://ieefa.org/resources/south-koreas-lng-overbuild.
[45] Shapiro A, Nemirovski A. On complexity of stochastic programming problems. Contin optim: Curr trends mod appl 2005:111–46.
[46] Yang Y, Luo Z, Yang Y, Wang D. Multistage adaptive distributionally robust optimization for the medical supplies distribution problem with uncertain demand in humanitarian aid. Comput Oper Res 2024;167:106631.
[47] Goerigk M, Kurtz J. Data-driven robust optimization using deep neural networks. Comput Oper Res 2023;151:106087.
[48] Ministry of Trade, Industry, and Energy, Ministry of Trade, Industry, and Energy. The 15th long-term natural gas supply and demand plan (2023–2036), Tech. rep., Sejong, South Korea, in Korean (Ministry of Trade, Industry and Energy; 2023.
[49] Kim J, Kim J. Optimal portfolio for LNG importation in korea using a two-step portfolio model and a fuzzy analytic hierarchy process. Energies 2018;11:3049.
[50] Ryu S. Global LNG outlook 2024-2h, Tech. rep., Daegu, South Korea, in Korean (Research Institute of Economics and Management, Korea Gas Corporation; 2025.
[51] Shin Y, Lee YP. Design of a boil-off natural gas reliquefaction control system for LNG carriers. Appl Energy 2009;86(1):37–44.
[52] Park C, Song K, Lee S, Lim Y, Han C. Retrofit design of a boil-off gas handling process in liquefied natural gas receiving terminals. Energy 2012;44(1):69–78.
[53] Feizabadi J. Machine learning demand forecasting and supply chain performance. Int J Logist Res Appl 2022;25(2):119–42.
[54] Tang CS. Perspectives in supply chain risk management. Int J Prod Econ 2006;103(2):451–88.
[55] Kleindorfer PR, Saad GH. Managing disruption risks in supply chains. Prod Oper Manag 2005;14(1):53–68.