Supply Chain Management and Economic Valuation of Real Options in the Natural Gas and Liquefied Natural Gas Industry

Mulan Xiaofeng Wang

Submitted to the Tepper School of Business
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

at

Carnegie Mellon University

Dissertation Committee:

Professor Nicola Secomandi (Chair)

Professor Alan Scheller-Wolf

Professor Sunder Kekre

Professor Ignacio Grossmann (outside reader)

Professor Duane Seppi (outside reader)

April 22, 2008

Abstract

My dissertation concentrates on several aspects of supply chain management and economic valuation of real options in the natural gas and liquefied natural gas (LNG) industry, including gas pipeline transportation, ocean LNG shipping logistics, and downstream storage.

Chapter 1 briefly introduces the natural gas and LNG industries, and the topics studied in this thesis.

Chapter 2 studies how to value U.S. natural gas pipeline network transport contracts as real options. It is common for natural gas shippers to value and manage contracts by simple adaptations of financial spread option formulas that do not fully account for the implications of the capacity limits and the network structure that distinguish these contracts. In contrast, we show that these operational features can be fully captured and integrated with financial considerations in a fairly easy and managerially significant manner by a model that combines linear programming and simulation. We derive pathwise estimators for the so called deltas and structurally characterize them. We interpret them in a novel fashion as discounted expectations, under a specific weighing distribution, of the amounts of natural gas to be procured/marketed when optimally using pipeline capacity. Based on the actual prices of traded natural gas futures and basis swaps, we show that an enhanced version of the common approach employed in practice can significantly underestimate the true value of natural gas pipeline network capacity. Our model also exhibits promising financial (delta) hedging performance. Thus, this model emerges as an easy to use and useful tool that natural gas shippers can employ to support their valuation and delta hedging decisions concerning natural gas pipeline network transport capacity contracts. Moreover, the insights that follow from our data analysis have broader significance and implications in terms of the management of real options beyond our specific application.

Motivated by current developments in the LNG industry, Chapter 3 studies the operations of LNG supply chains facing both supply and price risk. To model the supply uncertainty, we employ a closed-queuing-network (CQN) model to represent upstream LNG production and shipping, via special ocean-going tankers, to a downstream re-gasification facility in the U.S, which sells natural gas into the wholesale spot market. The CQN shipping model analytically generates the unloaded amount probability distribution. Price uncertainty is captured by the spot price, which experiences both volatility and significant seasonality, i.e., higher prices in winter. We use a trinomial lattice to model the price uncertainty, and calibrate to the extended forward curves. Taking the outputs from the CQN model and the spot price model as stochastic inputs, we formulate a real option inventory-release model to study the benefit

of optimally managing a downstream LNG storage facility. This allows characterization of the structure of the optimal inventory management policy. An interesting finding is that when it is optimal to sell, it is not necessarily optimal to sell the entire available inventory. The model can be used by LNG players to value and manage the real option to store LNG at a re-gasification facility, and is easy to be implemented. For example, this model is particularly useful to value leasing contracts for portions of the facility capacity. Real data is used to assess the value of the real option to store LNG at the downstream re-gasification facility, and, contrary to what has been claimed by some practitioners, we find that it has significant value (several million dollars).

Chapter 4 studies the importance of modeling the shipping variability when valuing and managing a downstream LNG storage facility. The shipping model presented in Chapter 3 uses a "rolling forward" method to generate the independent and identically distributed (i.i.d.) unloaded amount in each decision period. We study the merit of the i.i.d. assumption by using simulation and developing an upper bound. We show that the model, which uses the i.i.d. unloaded amount, provides a good estimation of the storage value, and yields a near optimal inventory control policy. We also test the performance of a model that uses constant throughput to determine the inventory release policy. This model performs worse than the model of Chapter 3 for storage valuation purposes, but can be used to suggest the optimal inventory control policy, especially when the ratio of flow rate to storage size is high, i.e., storage is scarce.

Chapter 5 summarizes the contributions of this thesis.

Acknowledgments

I started working on research topics in the natural gas industry four years ago. A great number of people supported me in assorted ways and helped keep my research adventure moving forward smoothly. It is a pleasure to express my gratitude to all of them here.

I am deeply grateful to Nicola (Nick) Secomandi, committee chairman, for his supervision, advice and guidance since the very early stages of this research. I remember the first day I met him. I planned to discuss the TA work with him, but I was totally enchanted by the amazing research world he opened to me. His passion about the OM/Finance interface infected me. I was so eager to do research with him. Since then, he has supported me in various ways. He frequently lent me quite a few books, suggested interesting papers to read, and gave me detailed guidance on how to learn finance, a completely new area for me. His rigorous attitude to science, constant encouragement, and the wonderfully useful discussions with him have been extremely important elements during my research growth in these four years, and inspired me to make remarkable progress in research.

Many thanks go in particular to Alan Scheller-Wolf and Sunder Kekre, my dissertation advisors, for their advice, supervision and crucial contribution to this thesis. I learned a lot from Alan about queuing theory, inventory control theory and other topics. Sunder always helped me to escape narrow views of research and guided me toward thinking in a much broader way. This has made my research work much more interesting and practically useful. Nick, Alan, and Sunder, I am grateful in every possible way and hope to keep up our collaboration in the future.

I am also thankful to Professor Duane Seppi and Professor Ignacio Grossmann. I am much indebted to them for the valuable advice and insightful discussion about my research and their precious time spent reading this thesis and giving me valuable comments.

I gratefully acknowledge the Center for Analytical Research in Technology (CART), Carnegie Mellon Tepper School of Business, for the financial support for my research work.

Special thanks go to Lawrance Rapp, Rosanne Christy, and Jackie Cavendish for their professional, prompt and considerate help, which made my life much easier allowing me to have more time for doing research. I also benefited a lot from discussions with all my fellow graduate students.

I gratefully thank my colleagues of BP Energy, Houston, TX. When I was an intern in their Quantitative Analysis group, they gave me many useful comments and suggestions about my research.

Words fail to express my appreciation to my parents and my little sister for their unconditional support over the years. Their love and persistent confidence in me encouraged me

to go through all the inevitable hard times, and stimulated me to work harder and achieve my research goals.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Natural Gas is a major principal source of energy. Since it is environmental clean, and there is abundance of natural gas in the world, many power plants use it as fuel. It plays an increasing important role in serving day-to-day energy needs.

The natural gas supply chain includes upstream production, in the United States (U.S.) or foreign countries, transportation, and storage. There are two types of transportation: natural gas pipeline and liquefied-natural-gas (LNG) ships. The pipeline-based transportation system consists of a complex network of pipelines, which can quickly and efficiently transport natural gas from supply areas to demand areas. However, it is impractical to use pipelines across oceans. LNG is natural gas cooled to liquid state. Liquefaction reduces the volume of natural gas by a factor of more than 600, thus makes storage and long distance ocean shipping practical. After being loaded to special ships, LNG is shipped to downstream storage facilities, where it is regasified and sold into the wholesale market through pipeline delivery. This thesis studies several aspects of supply chain management and economic valuation of real options in the natural gas and LNG industries, including gas pipeline transportation, ocean LNG shipping logistics, and LNG downstream storage.

Natural gas prices in the U.S. are set by market forces. Buying and selling of natural gas by market players drives the movement of natural gas prices, which can change often, and sometimes drastically over time. Thus, natural gas prices in the U.S. are very volatile. There are two distinct markets for natural gas: the spot market, and the futures market. The spot market entails transactions with daily delivery of natural gas. The futures market consists of trades with delivery of natural gas at least one month, and up to 72 months, in the future. Physical delivery under a futures contract is rare. Thus, futures trading is mainly a financial, rather than physical, trading activity. Market players also trade other natural gas based financial derivatives, such as options and basis swaps. The New York Mercantile Exchange (NYMEX) and the IntercontinentalExchange (ICE) are the main exchanges for

the trading of these contracts.

The main players in the natural gas industry are producers, interstate pipeline companies (IPCs), local distribution companies (LDCs), marketers, and end users. Natural gas wholesale markets in the U.S. have been deregulated since the mid-1980s. The pipeline system is owned and operated by pipeline companies. Interstate pipelines are regulated by the Federal Energy Regulatory Commission (FERC). FERC requires that pipeline companies make their transport capacity available to shippers in a non-discriminatory basis. Shippers must contract with pipeline companies to use their transport capacity. Transport contracts can be viewed as real options: when the natural gas price at the delivery market is higher than the price at the receipt market net of transportation cost, shippers use transport contracts to capture the positive price differences; otherwise, they do not use their transport contracts.

There are two types of transport contracts: firm and interruptible. Interruptible contracts are usually short-term contracts. Under this type of contract, the pipeline companies have no obligation to provide the transport service. Firm contracts give shippers the right to transport up to a specified amount of natural gas between specified locations, and require the pipeline companies to guarantee service. Firm contracts have an option-like structure: shippers pay a premium (demand charge) to pipeline companies to reserve transport capacity and an execution fee (strike price or commodity charge) to physically ship the gas. The focus of Chapter 2 in this thesis is on firm contracts.

Most transport contracts have a network structure that includes multiple receipt/delivery points, that is, shippers can ship gas from several receipt markets to several delivery markets under the same contract. Some contracts have dedicated capacity, that is, each receipt or delivery point has a pre-specified maximum capacity. Shippers cannot transport more than this amount of gas from (or to) each point. Flexible receipt/delivery contracts allow shippers to choose where they want to buy/sell natural gas, so long as the total quantity shipped does not exceed the contract capacity.

It is common for natural gas shippers to value transport contracts by simple adaptations of financial spread option formulas, which do not fully account for the implications of the capacity limits and the network structure that distinguish these contracts. Chapter 2 studies the real option valuation problem of transport capacity by taking the network structure and capacity flexibility of transport contracts fully into account.

Motivated by current developments in the LNG industry worldwide, Chapter 3 studies the operations of LNG supply chains facing both supply and price risk. In 2006, U.S. gas imports accounted for 15% of total gas consumption: 2% are LNG imports and 13% are pipeline imports. According to the Energy Information Administration (EIA), total U.S. gas imports will rise to 28% of total gas consumption by 2025, LNG imports will increase to

20%, and pipeline imports will decrease to 8%. Thus, LNG is expected to play an increasingly important role in the U.S. This expansion requires the development of more import facilities in the U.S. There are currently four existing conventional LNG import terminals, and 40 proposed or under-construction projects in the U.S. The existing four terminals are also expanding their storage and regasification capacities. This growing industry needs a model to support decisions on developing, valuing, and managing LNG import terminals, which can be interpreted as real options on the price of natural gas.

Chapter 3 studies the integrated LNG supply chain with a focus on the valuation of downstream storage and regasification terminals. Chapter 4 studies the importance of modeling shipping variability when valuing and managing downstream LNG storage.

Chapter 2

Computing and Delta Hedging the Economic Value of Natural Gas Pipeline Network Capacity

2.1 Introduction

Natural gas is a major energy source in the United States (U.S.) and other industrialized countries; in the U.S. it accounts for 24% of total energy consumption. Furthermore, the Energy Information Administration (EIA 2006 [19]) has projected that gas consumption of the U.S. will steadily grow by 1% per year from 2006 to 2030. In North America natural gas is traded on both spot and forward markets at different geographical locations. Gas Daily, a widely circulated industry newsletter, includes 84 pricing points. These locational markets are connected by a web of about 160 interstate pipelines.

The New York Mercantile Exchange (NYMEX) and the IntercontinentalExchange (ICE) trade financial contracts associated with about 40 locational markets. These contracts include futures with physical delivery at Henry Hub, Louisiana, basis swaps (forward financial contracts on price differences between Henry Hub and other markets), and put and call options on NYMEX futures. There are also over the counter (OTC) markets. Together NYMEX, ICE, and these OTC markets provide market participants with a high level of price transparency.

While natural gas wholesale markets in the U.S. are unregulated, interstate pipeline companies are regulated entities. They act as common carriers that do not own the gas they transport. Shippers, that is, those who own the gas being transported by pipelines, must contract with pipeline companies for portions of their transport capacity to receive service. Pipelines sell their capacity through sealed bid auctions run on their web sites. Interstate

Table 2.1: Points and quantities of a network transport contract on the Transco pipeline (Source: Transco web site).

Point	Type	Capacity (MMBtu/day)
Zone 3	Delivery	1,816
Zone 4	Delivery	111,366
Zone 1	Receipt	18,932
Zone 2	Receipt	27,841
Zone 3	Receipt	21,160
Zone 3	Receipt	43,433
Zone 3	Receipt	1,816

pipeline *minimum* and *maximum* contract rates (prices) are regulated, and shippers need to form their own valuations of pipeline capacity when they bid for securing it. Pipelines face a similar problem when determining a minimum acceptable price for their capacity.

Pipeline companies sell two basic types of transport contracts: *firm* and *interruptible*. Interruptible contracts are generally short term (1 month or less) best effort contracts, that is, pipelines are not obligated to provide interruptible transportation services to shippers. In contrast, firm contracts are long(er) term guaranteed contracts that give shippers the right to transport up to a given quantity of natural gas during each period in their term. Firm contracts have an interesting option-like structure: shippers pay a per unit premium (demand charge) to reserve transport capacity and a per unit execution fee (commodity charge) to use it. A large fraction of pipeline capacity in the U.S. is sold in advance on a firm basis.

As discussed by Eydeland and Wolyniec (2003 [21]), shippers can value pipeline capacity as a real option on natural gas prices at different locations connected by a pipeline. In particular, pipeline capacity between *two* locations can be valued as a spread option on natural gas prices at these locations. Their argument is that shippers employ pipeline capacity to support the following trade: they purchase natural gas at a receipt market and inject it into the pipeline, which in turn transports and delivers it to the delivery market, where shippers finally sell the gas on the spot market. (To be precise, this argument applies when the shipper is a merchant but can be shown to hold when the shipper is a natural gas producer, industrial consumer, or local distribution company; see Secomandi 2007 [52].)

Secomandi (2007 [52]) provides some empirical evidence that shippers do in fact use spread option pricing methods to value point-to-point pipeline transport capacity. But not all transport contracts have this simple structure; they often feature a *network* structure that

includes multiple receipt/delivery points. That is, shippers can ship gas from several receipt markets to several delivery markets under the same contract. Table 2.1 shows the points and capacities of a real network transport contract on the Transco pipeline (the meaning of the point capacities is explained in more detail in §2.2). Conversations with practitioners reveal that simple adaptations of the spread option pricing approach are widespread in practice for the valuation of network contracts. However, their network structure and capacity limits imply that in this case the spread option valuation logic may no longer apply in exact terms. Thus, in general, these spread option based methods are heuristics that yield suboptimal valuations of network contracts for pipeline capacity.

Objectives and contributions. To the best of our knowledge, the effectiveness of spread option based heuristic models is not currently known. Our first objective is to develop an exact model to compute the economic value of network pipeline capacity, and compare its performance against that of heuristic spread option based models. Our second objective is to assess the practical viability of pricing this capacity by risk neutral valuation techniques, that is, by assuming that financial replication of the economic value of this capacity by dynamic trading of futures (delta hedging) is possible in practice. Our third objective is to study how the contract value changes in price-related parameters, including initial prices, volatilities, and correlations.

We contribute to the literature as follows. We develop and analyze a linear programming and simulation (LPS) based model that, aside from simulation error, exactly values natural gas pipeline network transport capacity. We establish a parity type decomposition of the value of this capacity that extends the classical put-call parity result for European call and put options (Hull 2000 [30]). We derive two spread option based models, respectively based on linear and convex programming, which yield a lower bound (LB) and an upper bound (UB) on the exact value of capacity. Our LB model appears to be an enhanced version of spread option based models employed in practice, and we take it as a proxy for current practice. Our UB provides an additional benchmark for our LPS model.

Using an Ito based representation of the evolution of futures prices at different geographical locations, we apply the pathwise derivative computation method (Glasserman 2004 [25]) to obtain unbiased estimators of the so called deltas, that is, the number of futures contracts to be dynamically traded to replicate the economic value of network pipeline capacity (Hull 2000 [30]). In addition, we establish structural properties of the deltas and interpret them as discounted expectations, under a specific weighing distribution, of the amount of natural gas procured/marketed when optimally using pipeline capacity.

We also conduct an analysis of the effectiveness of our LPS model relative to the L/UB models based on real natural gas prices. We show that the LB model can significantly

undervalue network transport capacity, by 5-25% in different cases, relative to our LPS model. The UB model largely overvalues this capacity, which further indicates that it seems difficult for spread option based models to generate accurate valuations of this capacity. Perfect delta hedging can capture the exact value change along the time. We test the delta hedging performance of our LPS model on real data and measure the discrepancy of estimate value at time 0 and total realized value at time T. If the contract value is perfectly delta hedged, the discrepany is zero. Our numerical results show that our LPS model obtains encouraging results.

In addition to the deltas, we are also interested in the sensitivity analysis with respect to the other price parameters. We use the central difference method to compute the other derivatives of price parameters, and numerical examples to illustrate how the contract values change with the initial futures prices, volatilities, correlations, and time to maturity.

Relevance. While practicing managers in the energy and commodity industries have embraced real option thinking and tools, application of these ideas tends to overlook important operational aspects of the real assets being managed. As previously discussed, in our setting this amounts to using spread option based models for the valuation of network transport capacity. Our work brings to light the importance of faithfully capturing the operational aspects of these assets when valuing them. Moreover, conversations with practitioners reveal that our formal interpretation of the deltas is consistent with how these quantities tend to be used in practice, that is, as "projections" for how much to buy/sell when a contract is used. Thus, this interpretation provides theoretical support for practice.

In particular, our LPS model is a tool that natural gas traders and risk managers of firms engaged in the shipping of natural gas can employ to value pipeline network capacity and support related trading and hedging activities. These firms include producers, merchants, industrial consumers, utilities, and local distribution companies. Our data driven assessment of model performance, both in terms of valuation and financial hedging abilities, the simplicity of our model, and its fast computation time make our LPS model a good candidate for practical implementation and adoption.

While our models and insights are specific to one type of real option in the natural gas industry, they have relevance beyond our specific application. Our models can be extended to value and hedge other real options in the energy and commodity industries associated with the management of refining, processing, shipping, transportation, and distribution capacity for various commodities, such as, oil, coal, and biofuels. We expect our insights to remain pertinent in these settings.

Our sensitivity analysis studies the movement of the contract value with respect to the price parameters. The analysis is interesting and useful for practitioners. Shippers can use

Table 2.2: Relationship of the LPS model to the existing models for exotic options' valuation.

Case	Receipt Set Cardinality	Delivery Set Cardinality	Receipt Set Capacity	Delivery Set Capacity	Model
1	1	$\geqslant 1$	Dedicated	Dedicated	Spread option
2	$\geqslant 1$	1	Dedicated	Dedicated	Spread option
3	> 1	> 1	Flexible	Flexible	Rainbow option
4	> 1	> 1	Dedicated	Dedicated	LPS
5	> 1	> 1	Flexible	Dedicated	LPS
6	> 1	> 1	Dedicated	Flexible	LPS

this information to support their trading activities in financial markets.

Novelty. Our work contributes to the real option literature that deals with applications in the energy and commodity industries. Geman (2005 [24]) provides a recent introduction to this field. Schwartz (1997 [50]) and Seppi (2002 [54]) review the stochastic processes typically used in this literature to model the evolution of spot prices. Smith and McCardle (1998 [57], 1999 [58]) discuss the valuation of oil and gas investments by focusing on production activities. Enders et al. (2008 [16]) value extraction and technology scaling options in natural gas production. Kamat and Oren (2002 [34]) and Baldick et al. (2006 [2]) value interruptible electricity contracts. Tseng and Barz (2002 [62]) and Tseng and Lin (2007 [63]) consider the valuation of power generation assets. Jaillet et al. (2004 [31]) and Keppo (2004 [36]) analyze the valuation of electricity and natural gas swing contracts. Wang et al. (2007 [64]) and Secomandi (2007 [53]) value liquefied natural gas and commodity storage assets, respectively. Caldentey et al. (2007 [8]) study the long-term operation of an underground copper mining project. Hahn and Dyer (2008 [26]) value oil and gas switching options. Martínez-de-Albéniz and Simón (2007 [42]) study the optimal trading of a commodity between two geographical locations when the trader has market power. In contrast to these authors, we study a different problem, the valuation of natural gas pipeline network transport capacity, which does not appear to have been studied in the literature.

Deng et al. (2001 [13]), Eydeland and Wolyniec (2003 [21]), and Secomandi (2007 [52]) consider simplified versions of the problem studied here, that is, the valuation of point-to-point electricity transmission and natural gas transportation assets. Given the point-to-point nature of these assets, their valuation through spread options is appropriate. In this paper we study the *network* version of this problem, which is significantly richer than its point-to-point counterpart and requires a different solution approach for accurate valuation purposes (our LPS model).

Our contract valuation problem can be interpreted as the valuation of an exotic option.

Special cases of this problem, characterized by different flexibilities in the usage of the capacity of different receipt and delivery points (see §2.2 for details), can be addressed using existing models available in the financial engineering literature (see, e.g., Hull 2000 [30, Chapter 18]). Table 2.2 relates our LPS model to this literature by varying the cardinalities and the capacity flexibilities of the receipt and delivery point sets. If a contract has either one receipt point or one delivery point and both the receipt and delivery point sets have dedicated capacities (cases 1-2), then the problem reduces to the valuation of a collection of spread options (see Carmona and Durrleman 2003 [9] for a review of spread option valuation models). If the receipt or delivery point sets have flexible capacity (case 3), then our contract valuation problem is equivalent to the valuation of a rainbow option (Stulz 1982 [60], Johnson 1987 [33], and Boyle and Tse 1990 [5] deal with the valuation of such options). Otherwise (cases 4-6), the existing models do not capture the entire richness of our problem, and our LPS model is needed to solve it exactly. Obviously, our model also applies in cases 1-3.

No exact closed form formulas are available for the valuation of spread and rainbow options under the typical models used to represent the evolution of equity or commodity prices. Thus, simulation must be used to value these options exactly, but there exist closed form approximate formulas to price these options (see the previously stated references). While we also use simulation and closed form spread option approximation formulas in this paper, a novel feature of our work is combining these techniques with linear/convex optimization methods.

Our work is distinct from most of the extant real option literature in that we assess both the valuation *and* delta hedging performance of our LPS model based on real data. The test of the delta hedging ability of valuation models on real data is not typically carried out in this literature. Our computation of the deltas is based on applying the pathwise method discussed by Broadie and Glasserman (1996 [7]) and Glasserman (2004 [25]). A significant element of novelty in our application of this technique is that the terminal payoff of our real option is given by the solution of a linear program, rather than being a closed form expression, as is typical in the valuation of equity options (Broadie and Detemple 2004 [6]). By invoking envelope theorems (Milgrom and Segal 2002 [43]), we can apply the pathwise method to our setting. Moreover, our interpretation of the delta hedges as special discounted expectations of the quantities optimally procured/marketed at contract expiration leverages, in a specific and novel manner, our pathwise delta expressions.

Our numerical examples show that the contract values are convex, decreasing first, and then increasing, in price volatility when the relevant prices are highly positively correlated. This "V" shape behavior is uncommon for single-asset options, but has been already observed for spread options (Eydeland and Wolyniec 2003 [21], pp.344-345). Our sensitivity analysis

shows that this behavior continues to occur for network-type options.

Organization. The remainder of this chapter is organized as follows. We present our LPS model in §2.2 and analyze it in §2.3. We introduce our L/UB models in §2.4. We discuss delta hedging related aspects in §2.5, and discuss the central difference method to compute the other Greeks in §2.6. We empirically assess the performance of our models and conduct our sensitivity analysis in §2.7, and briefly conclude in §2.8. An appendix contains supportive materials.

2.2 Model

Figure 2.1(a) conceptually illustrates a natural gas pipeline network. Gas flows from receipt points, through connecting points, to delivery points. Pipeline managers market their capacity in the form of contracts that specify a collection of time periods (term), two sets of receipt and delivery points, the capacity limits at each of these points, and the set of links that connect the receipt points to the delivery points.

The point capacities indicate how much natural gas a shipper owning such a contract can inject/withdraw at each receipt/delivery point in each time period in the contract term. These contracts do not specify link capacities, which are instead implicitly expressed by the point capacities. The contract capacity is equal to the sum of the capacities of the receipt or the delivery points. Thus, the logistical structure of a network contract can be represented as a bipartite graph with node capacities. We denote the sets of receipt and delivery points (nodes) by $\mathcal{R}$ and $\mathcal{D}$, respectively, and let m and n be their cardinalities. Figure 2.1(b) illustrates three possible (m:n) contract networks: (1:n), (m:n), and (m:1).

Two typical contract terms are the November-March (heating season) and April-October terms, but multiyear contracts are also common. A contract is really a collection of subcontracts, one for each period in the contract term, because during each such period shippers have the option to ship natural gas up to the maximum contract capacity, that is, capacity unutilized in a given period cannot be utilized at a later time. Thus, without loss of generality, we only consider contracts with a single period term, which we denote by T. By a slight abuse of notation we also denote by T the time that corresponds to the beginning of this period.

The receipt and delivery points of a network contract are associated with natural gas market hubs. During time period T, the shipper owning this contract purchases natural gas at one or more receipt points and ships it to one or more delivery points, where the delivered natural gas is sold. Purchases and sales occur during the same time period because natural gas is received and delivered simultaneously by the pipeline (but, clearly, the received

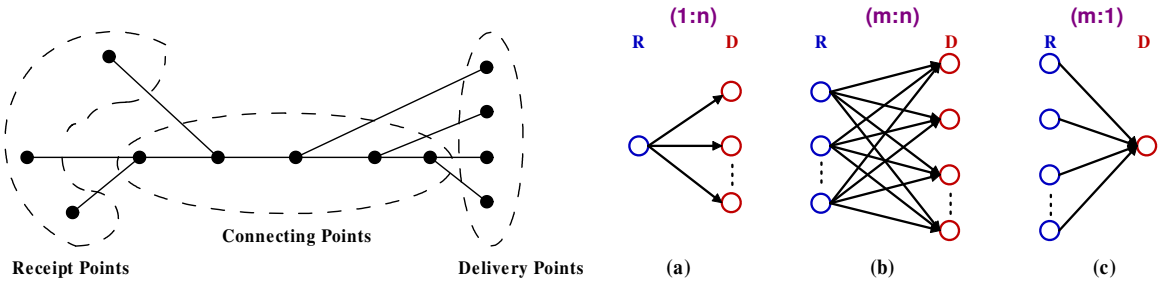

(a) Pipeline Network (b) Contract Network

Figure 2.1: Conceptual representations of natural gas pipeline and contract networks.

natural gas is not the same delivered natural gas). The problem studied in this paper is that of computing at the current time, which we denote 0, the economic value of a given network contract for usage of pipeline transport capacity during time period T ($\geqslant 0$).

This problem can be tackled via risk neutral valuation techniques (see, e.g., Luenberger 1998 [39, Chapters 8-9] and Duffie 2001 [15, Chapters 2 and 6]). This requires that a futures market exists at each point of a given network transport contract. In the U.S., the main financial contracts relevant for valuation purposes are NYMEX natural gas futures, whose delivery point is Henry Hub, Louisiana. This is a very liquid market with up to 72 different monthly maturities. In addition, NYMEX and ICE trade basis swaps, which are locational price differences relative to Henry Hub. Basis swaps are purely financial contracts that do not entail physical delivery.

While there exist fewer basis swaps than physical markets, a basis swap typically "covers" more than one market hub. This is quite natural, as price movements at more than one market hub are closely correlated, so that given the costs of maintaining a financial contract (basis swap) fewer financial contracts have been developed than physical markets. For example, the Houston Ship Channel basis swap is clearly associated with the Houston Ship Channel physical market, but this basis swap can also be associated with the nearby Katy Hub, which does not feature a separate basis swap. Also, while futures exist only at Henry Hub, by definition of basis swap it is clear that the sum of the Henry Hub futures price for a given maturity and the basis swap price for the same maturity yields a futures price for the basis swap location. Hence, valuation by risk neutral methods is possible even for contracts that involve a large number of markets.

Natural gas futures contracts are settled three business days prior to the beginning of a given month. Time T corresponds to this time. Thus, we assume that this is when the shipper decides how much natural gas to ship during time period T, that is, the ensuing month. This also occurs in practice when shippers "nominate" their monthly shipping decisions to

pipelines during the so called *bid week*, which is the week prior to each shipping month (Eydeland and Wolyniec 2003 [21, Chapter 1]). However, shippers can modify their monthly nominations within the shipping month, e.g., in response to changes in the spot prices relative to the settled futures prices. We do not model these daily nomination updates, but our model can be modified in a straightforward manner to capture their effect on the contract value by dividing monthly time period T into weekly or daily subperiods. Thus, replication of the contract cash flows during time period T can be performed by trading futures contracts at each receipt/delivery point. If T were to be less than one month, this replication could be performed in practice using balance of the month/week or Gas Daily options (see Eydeland and Wolyniec 2003 [21, Chapter 4]).

We introduce some additional notation to describe our LPS model. The capacities of each receipt point $i \in \mathcal{R}$ and delivery point $j \in \mathcal{D}$ are denoted C_i and C_j, respectively, and are measured in MMBtu per month. The contract capacity is then $C := \sum_{i\in\mathcal{R}} C_i \equiv \sum_{j\in\mathcal{D}} C_j$. The per unit usage charge to ship natural gas during time period T on link i-j, $i \in \mathcal{R}$, $j \in \mathcal{D}$, is denoted K_{ij}; this is the commodity rate of link i-j and is measured in \$/MMBtu. The fuel coefficient associated with shipping natural gas on link i-j during time period T is $\phi_{ij} \in [0, 1]$; this coefficient is used to model the fuel required by the compressor stations to ship one unit of natural gas on this link, which is $\phi_{ij}/(1 - \phi_{ij})$, and must be procured by the shipper. The futures prices at time $t \in [0, T]$ with maturity at time T for receipt and delivery points $i \in \mathcal{R}$ and $j \in \mathcal{D}$ are $F_i(t, T)$ and $G_j(t, T)$, respectively. (If $t = T$ then these are spot prices.) The risk free interest rate is r; we employ continuous compounding in this paper so that the risk free discount factor from time $t \geqslant 0$ back to time 0 is $\exp(-rt)$.

Let $v(T)$ denote the value of the contract cash flows at contract execution time T for a given realization of prices $F_i(T, T)$ and $G_j(T, T)$, $\forall i \in \mathcal{R}$, $j \in \mathcal{D}$ ($v(T)$ clearly depends on these prices but for notational simplicity this dependency is omitted from our notation). This quantity is the optimal value of the following linear program P, where sets $\mathcal{D}(i)$ and $\mathcal{R}(j)$, respectively, include the delivery and receipt points connected to receipt point $i \in \mathcal{R}$ and delivery point $j \in \mathcal{D}$:

$$\text{P:} \quad v(T) := \max_{x} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij} \right] x_{ij} \tag{2.1}$$

$$\text{s.t.} \quad \sum_{j\in\mathcal{D}(i)} x_{ij} \leqslant C_i, \ \forall i \in \mathcal{R} \tag{2.2}$$

$$\sum_{i\in\mathcal{R}(j)} x_{ij} \leqslant C_j, \ \forall j \in \mathcal{D} \tag{2.3}$$

$$x_{ij} \geqslant 0, \ \forall i \in \mathcal{R}, j \in \mathcal{D}(i). \tag{2.4}$$

In model P we assume that all cash flows incurred during time period T are accounted for at

time T. As of time 0, the optimal objective function value of model P is a random variable, denoted $\tilde{v}(T)$, which depends on the uncertain prices at the receipt and delivery points that will prevail at time T. The contract value at time 0, denoted $V(0,T)$, is the expected value computed under the risk neutral measure of random variable $\tilde{v}(T)$ discounted by the risk free factor:

$$V(0,T) := e^{-rT} E_0[\tilde{v}(T)].$$

Here E_0 denotes the time 0 conditional expectation with respect to the joint distribution of random vector $(\tilde{F}_i(T,T), \tilde{G}_j(T,T), i \in \mathcal{R}, j \in \mathcal{D})$ under the risk neutral measure given price vector $(F_i(0,T), G_j(0,T), i \in \mathcal{R}, j \in \mathcal{D})$. This measure exists under standard assumptions that we suppose to hold here (Hull 2000 [30] and Luenberger 1998 [39]). These assumptions are fairly realistic in our setting given the existence of traded futures contracts at different geographical locations in North America, as previously discussed. We compute $V(0,T)$ by simulating time T futures prices' realizations, optimally solving linear program P for each such realization to obtain $v(T)$, averaging these values, and discounting back to time 0.

Model P assumes that the point capacities are dedicated. Some contracts in practice have a more flexible structure. Contracts with flexible capacities allow shippers to combine the capacities of different receipt/delivery points for usage at a single such point. Three broad classes of flexible capacity contracts include those with capacity flexibility associated with either their receipt or delivery points, but not both, and those with fully flexible receipt and delivery points (of course one may think of even more cases than these, in which only a subset of points have flexible capacities). For brevity, we do not provide the modified formulations of model P for these three cases.

2.3 Analysis

In this section we discuss properties of the optimal solution of model P and decompose the contract value as the sum of its intrinsic and extrinsic values (defined below).

Optimal solution of model P. An optimal solution to model P can be computed by linear programming. Since the cardinalities of the receipt and delivery point sets are small (less than 10 in practice), this can be done very efficiently, even if large cardinality LP is efficient. In fact one may question why linear programming is needed for this purpose; in other words, one may wonder if the following simple *greedy algorithm* (GA) would be sufficient to optimally solve model P.

Given a sample realization of futures prices at time T:

Step 1. Set $v^{GA}(T) \leftarrow 0$, $\mathrm{RC}_i \leftarrow C_i$, $\forall i \in \mathcal{R}$, and $\mathrm{DC}_j \leftarrow C_j$, $\forall j \in \mathcal{D}$.

Step 2. Set price spread coefficients $S_{ij}(T,T) \leftarrow G_j(T,T) - F_i(T,T)/(1-\phi_{ij}) - K_{ij}$, $\forall i \in \mathcal{R}$,

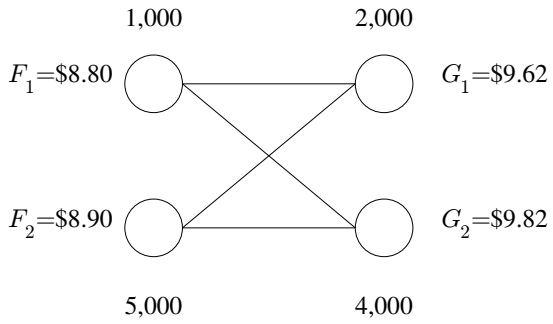

Figure 2.2: Contract network for Example 1.

$j \in \mathcal{D}(i)$, sort them in decreasing order, and store them in a stack.
Step 3. If the stack is empty, stop and return $v^{GA}(T)$. Otherwise, remove the first element $S_{ij}(T,T)$ from the stack. Set $x_{ij} \leftarrow \min\{\mathrm{RC}_i, \mathrm{DC}_j\}$ and $v^{GA}(T) \leftarrow v^{GA}(T) + S_{ij}(T,T)x_{ij}$.
Step 4. Set $\mathrm{RC}_i \leftarrow \mathrm{RC}_i - x_{ij}$ and $\mathrm{DC}_j \leftarrow \mathrm{DC}_j - x_{ij}$. Go to Step 3.

Example 1 shows that GA can fail to yield an optimal solution to P, even in rather simple settings.

Example 1 (GA suboptimality). Consider the 2 receipt and 2 delivery point dedicated capacity contract displayed in Figure 2.2, where the numbers above/below each point are their capacities. Let the fuel rate of each link be zero. The time T prices are $F_1 = \$8.80$, $F_2 = \$8.90$, $G_1 = \$9.62$, and $G_2 = \$9.82$, with their MMBtu units and arguments removed for ease of notation (a simplification also made below with reference to other quantities of interest). Let the commodity rates be $K_{11} = \$0.01$ and $K_{12} = K_{21} = K_{22} = \$0.02$. The price spreads net of the commodity rates are $S_{11} = \$0.81$, $S_{12} = \$1.00$, $S_{21} = \$0.7$, and $S_{22} = \$0.9$. The GA solution is $x_{11}^{GA} = 0$, $x_{12}^{GA} = 1,000$, $x_{21}^{GA} = 2,000$, and $x_{22}^{GA} = 3,000$, and its value is $v^{GA}(T) = \$5,100$. An optimal solution to P is instead $x_{11}^{*} = 1,000$, $x_{12}^{*} = 0$, $x_{21}^{*} = 1,000$, and $x_{22}^{*} = 4,000$, and its value is $v(T) = \$5,110$. If K_{11} is set equal to the other commodity rates, that is, \$0.02, the GA solution is unchanged but is now also optimal (the previous optimal solution also remains so).

The cause of the suboptimality of GA in Example 1 resides in the point capacity constraints. In this example, link 1-2 has the highest spread ($S_{12} = \$1.00$), so one would be tempted to send as much gas as possible through this link. However, the capacity of receipt point 1 is 1,000. This means that if link 1-2 is fully utilized, then the only way to utilize the capacity of delivery point 1 is to ship gas from receipt point 2 to this delivery point, which implies utilizing the link with the lowest spread ($S_{21} = \$0.70$). Thus, it is optimal to avoid using the highest valued link 1-2, fully utilize the capacity of delivery point

1 by shipping gas along links 1-1 and 2-1, and fully utilize the capacity of delivery point 2 by shipping gas here from receipt point 2. More formally, starting from the GA solution, the contract network exhibits the feasible improving cycle ($i = 1$, $j = 1$, $i = 2$, $j = 2$, $i = 1$): the change in objective function value of sending one unit of flow through this cycle is $S_{11} - S_{21} + S_{22} - S_{12} = K_{21} - K_{11} + K_{12} - K_{22} = 0.01 > 0$, and the resulting solution satisfies all the capacity constraints. Hence, the GA solution can be made optimal by *not* shipping gas along the highest valued link. Thus, somewhat surprisingly, as a consequence of the point capacity constraints, it is optimal to ship gas along the least valued link but not along the highest valued link.

In Example 1, all the fuel rates are equal. When all the links also have identical commodity rates then GA optimally solves P. This is an instance of the following general result.

Proposition 1 (GA optimality). *GA optimally solves model P if all the fuel rates ϕ_{ij} and commodity rates K_{ij}, $\forall i \in \mathcal{R}$, $j \in \mathcal{D}(i)$, are equal to constants ϕ and K, respectively.*

To gain intuition into the nature of this result, it is useful to introduce purchase and sale decision variables $p_i = \sum_{j \in \mathcal{D}(i)} x_{ij}$ and $s_j = \sum_{i \in \mathcal{R}(j)} x_{ij}$ for receipt and delivery points i and j, respectively, and to define $F_i'(T,T) := F_i(T,T)/(1-\phi) - K$. In this case, model P can be equivalently reformulated as follows:

$$
\begin{aligned}
\text{P}'\text{:}\ v(T) \ :=\ & \max_{p,s} \sum_{j \in \mathcal{D}} G_j(T,T) s_j - \sum_{i \in \mathcal{R}} F_i'(T,T) p_i \\
\text{s.t.}\quad & 0 \leqslant p_i \leqslant C_i,\ \forall i \in \mathcal{R} \\
& 0 \leqslant s_j \leqslant C_j,\ \forall j \in \mathcal{D} \\
& s_j \leqslant \sum_{i \in \mathcal{R}(j)} p_i,\ \forall j \in \mathcal{D} \\
& p_i \leqslant \sum_{j \in \mathcal{D}(i)} s_j,\ \forall i \in \mathcal{R} \\
& \sum_{i \in \mathcal{R}} p_i = \sum_{j \in \mathcal{D}} s_j.
\end{aligned}
$$

It is now clear that model P$'$ can be solved to optimality in a sequential fashion as follows. Form two lists by ordering the selling prices in decreasing order and the purchasing prices in increasing order. Make *equal* sale and purchase profitable decisions, in the sense that the sale revenue exceeds the purchase cost, by repeatedly selecting the locations corresponding to the top elements of each of the two lists, setting their sale and purchase amounts equal to the minimum of their respective remaining point capacities, updating these capacities accordingly, and removing a selling/purchase price whenever the remaining capacity of its corresponding point becomes zero. This algorithm is evidently equivalent to GA.

Contract value decomposition. Analogous to the well known put-call parity result in the valuation of standard options (Hull 2000 [30]), we now establish an extended parity result for the valuation of a network contract. Given nonnegative and finite vector $x(0) := (x_{ij}(0), \forall i \in \mathcal{R}, j \in \mathcal{D}(i))$, where the suffix indicates that this vector is determined at time 0, we define

$$\begin{aligned}
V^1(0,T;x(0)) &:= \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} e^{-rT}\left[G_j(0,T) - \frac{F_i(0,T)}{1-\phi_{ij}} - K_{ij}\right] x_{ij}(0) \\
V^2(0,T;x(0)) &:= e^{-rT}E_0[\tilde{v}^2(T;x(0))] \\
v^2(T;x(0)) &:= \max_{x'(T)} \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij}\right] x'_{ij}(T) \qquad (2.5) \\
&\quad \text{s.t.} \sum_{j\in\mathcal{D}(i)} x_{ij}(0) + x'_{ij}(T) \leqslant C_i,\ \forall i \in \mathcal{R} \qquad (2.6) \\
&\quad \sum_{i\in\mathcal{R}(j)} x_{ij}(0) + x'_{ij}(T) \leqslant C_j,\ \forall j \in \mathcal{D} \qquad (2.7) \\
&\quad x'_{ij}(T) \geqslant -x_{ij}(0),\ \forall i \in \mathcal{R}, j \in \mathcal{D}(i). \qquad (2.8)
\end{aligned}$$

The vector $x'(T)$ in (2.5)-(2.8) can be interpreted as an update of the flow vector $x(0)$ at time T given the realization of prices at this time. Continuing the analogy with the well know put-call parity result, $V(0,T)$ corresponds to the value of a "call" option and $V^2(0,T;x(0))$ to that of a "put" option. Define $x_{ij}(T) := x_{ij}(0) + x'_{ij}(T)$, $\forall i \in \mathcal{R}$, $j \in \mathcal{D}(i)$. Replace $x'_{ij}(T)$ by $x_{ij}(T) - x_{ij}(0)$ in (2.5)-(2.8), so that the decision variables of this linear program become $x_{ij}(T)$. It is clear that these variables satisfy the constraints of model P. By defining $x^*(T)$ to be an element of

$$\arg\max_{x(T)} \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij}\right] [x_{ij}(T) - x_{ij}(0)] \text{ s.t. (2.6)-(2.8)},$$

and denoting by $\tilde{x}^*_{ij}(T)$ the random variable associated with one of its elements, we can write

$$\begin{aligned}
V^2(0,T;x(0)) &= e^{-rT}E_0\left\{\sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} \left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi_{ij}} - K_{ij}\right] [\tilde{x}^*_{ij}(T) - x_{ij}(0)]\right\} \\
&= e^{-rT}E_0\left\{\sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} \left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi_{ij}} - K_{ij}\right] \tilde{x}^*_{ij}(T)\right\} \\
&\quad - \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} e^{-rT}E_0\left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi_{ij}} - K_{ij}\right] x_{ij}(0) \\
&= V(0,T) - \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{D}(i)} e^{-rT}\left[G_j(0,T) - \frac{F_i(0,T)}{1-\phi_{ij}} - K_{ij}\right] x_{ij}(0) \\
&= V(0,T) - V^1(0,T;x(0)),
\end{aligned}$$

where the third equality follows since $F_i(0,T) = E_0[\tilde{F}_i(T,T)]$ and $G_j(0,T) = E_0[\tilde{G}_j(T,T)]$. Thus, the above analysis shows that the following result holds.

Proposition 2 (Network contract value parity). *Given* $0 \leqslant x(0) < \infty$ *(componentwise), it holds that* $V(0,T) = V^1(0,T;x(0)) + V^2(0,T;x(0))$.

Let $x^I(0)$ denote an optimal solution of the following intrinsic value model, where the adjective intrinsic indicates that the time 0 futures prices appear in the objective function of the following model, and we ignore price uncertainty:

$$V^I(0,T) := \max_{x(0)} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} e^{-rT} \left[G_j(0,T) - \frac{F_i(0,T)}{1-\phi} - K_{ij} \right] x_{ij}(0) \text{ s.t. (2.2)-(2.4).}$$

This allows us to decompose the contract value $V(0,T)$ into the sum of intrinsic value $V^I(0,T) \equiv V^1(0,T;x^I(0))$ and extrinsic value $V^E(0,T) := V^2(0,T;x^I(0))$: $V(0,T) = V^I(0,T) + V^E(0,T)$, where the adjective extrinsic indicates that this is the value that can be attributed to price uncertainty.

2.4 Bounds

In this section we derive lower and upper bounds on $V(0,T)$. The intrinsic value $V^I(0,T)$ is also a lower bound (LB) on $V(0,T)$, but we derive the following stronger LB denoted $V^{LB}(0,T)$:

$$\mathrm{P}^{LB}: \; V^{LB}(0,T) := \max_{x(0)} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} e^{-rT} E_0 \left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi} - K_{ij} \right]^+ x_{ij}(0) \text{ s.t. (2.2)-(2.4).}$$

That $V^{LB}(0,T) \leqslant V(0,T)$ follows from the observation that any optimal solution to P^{LB} is feasible to P over all sample path. The coefficients of the decision variables in the objective function of P^{LB} are spread options. These values can be computed by simulation or numerical methods, and good closed form approximation formulas, such as Kirk's (1995 [37]; see Appendix 2.9.1), are also available (Carmona and Durrleman 2003 [9]). Obviously if such approximations are used one also obtains an approximate LB.

We compare the two LBs, then have $V^{LB}(0,T) \geqslant V^{I}(0,T)$, because

$$\begin{aligned} V^{I}(0,T) &= \max_{x(0)} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} e^{-rT} \left[G_j(0,T) - \frac{F_i(0,T)}{1-\phi} - K_{ij} \right] x_{ij}(0) \\ &= \max_{x(0)} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} e^{-rT} \left[E_0[G_j(T,T)] - \frac{E_0[F_i(T,T)]}{1-\phi} - K_{ij} \right] x_{ij}(0) \\ &\leqslant \max_{x(0)} \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} e^{-rT} E_0 \left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi} - K_{ij} \right]^{+} x_{ij}(0) \\ &= V^{LB}(0,T) \end{aligned}$$

The first equality follows from the well known property that futures prices are martingales under the risk neutral measure (also used in establishing Proposition 2). Therefore, the LB model provides a tighter estimation than the intrinsic model.

We have anecdotal evidence that P^{LB} is related to spread option based models employed in practice. We conjecture that such models differ from P^{LB} in that practitioners solve it using a simple greedy search algorithm analogous to the one presented in §2.3. Thus, with dedicated capacity, such models provide a weaker LB on $V(0,T)$ than $V^{LB}(0,T)$ does.

We use Lagrangian duality applied to P to obtain our UB. Let λ_i and μ_j, respectively, be the Lagrange multipliers for each constraint (2.2) and (2.3) in P. The Lagrangian function for P is

$$\begin{aligned} L(x,\lambda,\mu;T) &= \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij} \right] x_{ij} - \sum_{i\in\mathcal{R}} \lambda_i \left[\sum_{j\in\mathcal{D}(i)} x_{ij} - C_i \right] \\ &\quad - \sum_{j\in\mathcal{D}} \mu_j \left[\sum_{i\in\mathcal{R}(j)} x_{ij} - C_j \right] \\ &= \sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij} - \lambda_i - \mu_j \right] x_{ij} + \sum_{i\in\mathcal{R}} \lambda_i C_i + \sum_{j\in\mathcal{D}} \mu_j C_j, \end{aligned}$$

where T in $L(x,\lambda,\mu;T)$ indicates that this problem is solved with knowledge of time T prices.

By Lagrangian duality, it holds that $v(T) = \min_{\lambda,\mu\geqslant 0} \max_{x\geqslant 0} L(x,\lambda,\mu;T)$. Each feasible decision variable x_{ij} in model P must satisfy the inequality $0 \leqslant x_{ij} \leqslant \min\{C_i, C_j\}$. Thus, to maximize $L(x,\lambda,\mu;T)$ over x, we set $x_{ij} = \min\{C_i, C_j\}$ if the coefficient of x_{ij} is positive and $x_{ij} = 0$ otherwise, which yields

$$v(T) = \min_{\lambda,\mu\geqslant 0} \underbrace{\sum_{i\in\mathcal{R}} \sum_{j\in\mathcal{D}(i)} \left[G_j(T,T) - \frac{F_i(T,T)}{1-\phi_{ij}} - K_{ij} - \lambda_i - \mu_j \right]^{+} \min\{C_i, C_j\} + \sum_{i\in\mathcal{R}} \lambda_i C_i + \sum_{j\in\mathcal{D}} \mu_j C_j}_{L'(\lambda,\mu;T)}.$$

Given vectors λ and μ, let $\tilde{L}'(\lambda, \mu; T)$ denote the random value of $L'(\lambda, \mu; T)$, e.g., as seen from time 0. By definition of $V(0,T)$, we obtain upper bound $V^{UB}(0,T)$ as follows:

$$V(0,T) = e^{-rT} E_0[\tilde{v}(T)] = e^{-rT} E_0 \left[\min_{\lambda,\mu \geqslant 0} \tilde{L}'(\lambda, \mu; T) \right] \leqslant \min_{\lambda,\mu \geqslant 0} e^{-rt} E_0[\tilde{L}'(\lambda, \mu; T)] =: V_{UB}(0,T).$$

Thus, we obtain the following result.

Proposition 3 (Bounds). *The following inequalities hold:*

$$V^I(0,T) \leqslant V^{LB}(0,T) \leqslant V(0,T) \leqslant V^{UB}(0,T).$$

Since it is easily verified that $L'(\lambda, \mu; T)$ is a jointly convex function of λ and μ, $V^{UB}(0,T)$ can be easily computed, provided that each of the following terms can be efficiently evaluated:

$$E_0 \left[\tilde{G}_j(T,T) - \frac{\tilde{F}_i(T,T)}{1-\phi_{ij}} - K_{ij} - \lambda_i - \mu_j \right]^+, \ \forall i \in \mathcal{R}, j \in \mathcal{D}(i).$$

These terms are spread options on futures prices' differences with a strike price (commodity rate) adjusted to account for the marginal values, as seen from time 0, of the capacity of receipt point i and delivery point j. As previously stated, there exist good closed form approximations for computing such option values. Similar to the LB case, if one employs such approximations one obtains an approximate UB.

2.5 Delta Hedging

It is well known that, in theory, financial replication of the value of a commodity derivative can be done by dynamically constructing portfolios of futures contracts (Black 1976 [4]). This is the celebrated delta hedging method, which requires the computation of the partial derivatives of the contract value with respect to changes in the current futures prices, that is, the deltas. In this section we show how delta hedging gives us a natural method to empirically test the valuations generated by our LPS model. We also discuss the computation of unbiased deltas, characterize, and interpret them by employing an Ito based model of futures price evolution. In practice, it is common to consider additional value sensitivities ("greeks") due to second order price effects or changes in other parameters associated with the model employed to describe the evolution of futures prices, e.g., their volatilities and correlations. For brevity, here we only focus on the deltas.

Delta hedging and LPS model testing. Since delta hedging is well known, we only provide a very brief discussion of how it applies in the context of this paper to make the exposition self contained. The receipt and delivery point deltas at time $t \in [0,T)$

are $\delta_i^F(t,T) := V_{F_i(t,T)}(t,T)$ and $\delta_j^G(t,T) := V_{G_j(t,T)}(t,T)$, where $V_\theta(t,T) := \partial V(t,T;\theta)/\partial\theta$ for $\theta \in \{F_i(t,T), G_j(t,T)\}$. Delta hedging consists of "shorting" the deltas at each time $t \in [0,T)$. In other words, the number of futures contracts to trade at time t at receipt and delivery points i and j are $M_i(t,T) := -\delta_i^F(t,T)$ and $N_j(t,T) := -\delta_j^G(t,T)$, respectively.

In theory, these futures positions must be continuously updated, in which case, letting $dF_i(s,T)$ and $dG_j(s,T)$ denote "generic" Ito based stochastic dynamics of prices $F_i(s,T)$ and $G_j(s,T)$, $\forall i \in \mathcal{R}$ and $j \in \mathcal{D}$, it is well known (see, e.g., Shreve 2004 [55]) that under some regularity conditions the change in the contract value between times 0 and T satisfies

$$\begin{aligned} V(T,T) - V(0,T) &= -\underbrace{\left[\sum_{i\in\mathcal{R}} \int_0^T M_i(s,T)dF_i(s,T) + \sum_{j\in\mathcal{D}} \int_0^T N_j(s,T)dG_j(s,T)\right]}_{\mathrm{CH}(0,T)} \\ \Rightarrow V(T,T) + \mathrm{CH}(0,T) &= V(0,T), \end{aligned}$$

where the term $\mathrm{CH}(0,T)$ is the cumulative delta hedging value between times 0 and T.

In practice, the delta positions can only be changed a finite number of times. Thus, suppose that updating of these positions occurs at the start of a finite number of discrete time intervals, each of length Δt, that partition the time interval $[0,T]$. Denote $\mathcal{T}(\Delta t)$ the set that includes the times corresponding to the *end* of each of these time intervals. Let $t \in \mathcal{T}(\Delta t)$ and define $\Delta F_i(t,T) := F_i(t,T) - F_i(t-\Delta t,T)$, $\Delta G_j(t,T) := G_j(t,T) - G_j(t-\Delta t,T)$, and

$$H(t,T) := \sum_{i\in\mathcal{R}} M_i(t-\Delta t,T)\Delta F_i(t,T) + \sum_{j\in\mathcal{D}} N_j(t-\Delta t,T)\Delta G_j(t,T).$$

In practice we can approximate $\mathrm{CH}(0,T)$ by $\widehat{\mathrm{CH}}(0,T) := \sum_{t\in\mathcal{T}(\Delta t)} H(t,T)$ to obtain

$$V(T,T) + \widehat{\mathrm{CH}}(0,T) \approx V(0,T). \tag{2.9}$$

Expression (2.9) is useful for empirically assessing the valuation yielded by our LPS model. Under a specific Ito based model of the futures prices' evolution, for example (2.10)-(2.14) below, our LPS model computes $V(0,T)$ at time 0. This value is obviously model based, that is, it depends on the probabilistic assumptions behind the futures price evolution model. However, we can compute both $V(T,T)$ and $\widehat{\mathrm{CH}}(0,T)$ on a sequence of real future prices observed between times 0 and T, for example, daily closing prices. These values are data rather than model based. (Even though the positions $M_i(\cdot,T)$ and $N_j(\cdot,T)$ that enter the computation of $\widehat{\mathrm{CH}}(0,T)$ are model based, the quantities $\Delta F_i(t,T)$ and $\Delta G_j(t,T)$ needed to compute $H(\cdot,T)$ are data based.) We refer to the ratio

$$\frac{V(T,T) + \widehat{\mathrm{CH}}(0,T) - V(0,T)}{V(0,T)}$$

as the discrepancy of our LPS model. Measuring this discrepancy on real price data provides us with an empirical method to assess the practical relevance of the valuations computed by our LPS model. While the logic underlying these arguments is well known to practitioners, as stated in §2.1, it is rarely employed in the academic literature to test the performance of a real option valuation model.

Delta computation, characterization, and interpretation. In the remainder of this section, we employ the following Ito based model of the evolution of futures prices at the receipt and delivery points under the *risk neutral* measure:

$$
\begin{aligned}
dF_i(t,T) &= F_i(t,T)\sigma_i^F(t,T)dW_i^F(t),\ \forall i \in \mathcal{R} && (2.10)\\
dW_i^F(t)dW_{i'}^F(t) &= \rho_{ii'}^{FF}(t)dt,\ \forall i,i' \in \mathcal{R} && (2.11)\\
dG_j(t,T) &= G_j(t,T)\sigma_j^G(t,T)dW_j^G(t),\ \forall j \in \mathcal{D} && (2.12)\\
dW_j^G(t)dW_{j'}^G(t) &= \rho_{jj'}^{GG}(t)dt,\ \forall j,j' \in \mathcal{D} && (2.13)\\
dW_i^F(t)dW_j^G(t) &= \rho_{ij}^{FG}(t)dt,\ \forall i \in \mathcal{R}, j \in \mathcal{D}. && (2.14)
\end{aligned}
$$

Here, $\sigma_i^F(t,T)$ is the volatility function (of time t) of the futures price at receipt point i with maturity at time T, $dW_i^F(t)$ is an increment to a standard Brownian motion, and $\rho_{ii'}^{FF}(t)$ is the instantaneous time t correlation between $dW_i^F(t)$ and $dW_{i'}^F(t)$. The notation for the futures price at delivery point j has an analogous interpretation, and $\rho_{ij}^{FG}(t)$ is the instantaneous time t correlation between $dW_i^F(t)$ and $dW_j^G(t)$, with $i \in \mathcal{R}$ and $j \in \mathcal{D}$.

Model (2.10)-(2.14) captures important models available in the literature. For example, it includes as a special case a generalized version of the celebrated Black (1976 [4]) model of commodity futures prices, whereby all the volatilities and instantaneous correlations are constants. Additionally, consider the model of Schwartz (1997 [50]), where the natural logarithm of the *spot* price at location $i \in \mathcal{R}$, denoted $X_i^F(t)$, follows the well known mean reverting process $dX_i^F(t) = \kappa_i^F[\xi_i^F - X_i^F(t)]dt + \sigma_i^F dW_i^F(t)$, with κ_i^F the speed of mean reversion, ξ_i^F the long-term level of $X_i^F(t)$, and σ_i^F its volatility. In this case, it holds that $\sigma_i^F(t,T) = \sigma_i^F \exp[-\kappa_i^F(T-t)]$. Given two distinct receipt points i and i', whose associated terms $dW_i^F(t)$ and $dW_{i'}^F(t)$ have constant instantaneous correlation $\rho_{ii'}^{FF}$, it is also easy to verify that $\rho_{ii'}^{FF}(t) = \rho_{ii'}^{FF}$. Analogous considerations apply to the delivery point futures prices and their correlations with the receipt point futures prices.

Below, we apply the pathwise method described by Glasserman (2004 [25, §7.2]) to obtain *unbiased* estimates of the deltas by simulation. Denote by $v(T;\theta)$ the dependence on scalar parameter θ of the optimal solution value to optimization problem P, and $\tilde{v}_\theta(T) := \partial\tilde{v}(T;\theta)/\partial\theta$ its pathwise partial derivative with respect to θ, that is, the random variable that can take values $dv(T;\theta)/d\theta$. A direct application of the pathwise derivative estimation

method would suggest that

$$V_\theta(t,T) = e^{-r(T-t)} E_t[\tilde{v}_\theta(T)], \forall t \in [0,T]. \tag{2.15}$$

But $v(T;\theta)$ is the optimal solution value to linear program P, a case not explicitly discussed in Glasserman (2004 [25, §7.2]). We can invoke envelope theorem results available in Milgrom and Segal (2002 [43], in particular their Corollary 4) to establish the validity of (2.15). This establishes part (a) of Proposition 4, while its part (b) is proved in Appendix 2.9.2.

Proposition 4 (Deltas). *(a) Under model (2.10)-(2.14), unbiased deltas at time $t \in [0,T)$ can be computed as follows:*

$$\delta_i^F(t,T) = -\frac{e^{-r(T-t)}}{F_i(t,T)(1-\phi)} E_t\left[\tilde{F}_i(T,T) \sum_{j\in\mathcal{D}(i)} \tilde{x}_{ij}^*(T)\right], \quad \forall i \in \mathcal{R} \tag{2.16}$$

$$\delta_j^G(t,T) = \frac{e^{-r(T-t)}}{G_j(t,T)} E_t\left[\tilde{G}_j(T,T) \sum_{i\in\mathcal{R}(j)} \tilde{x}_{ij}^*(T)\right], \quad \forall j \in \mathcal{D}. \tag{2.17}$$

(b) Moreover, the value $V(t,T)$ is convex and decreasing in each time t receipt point initial futures price, and convex and increasing in each time t delivery point initial futures price, or, equivalently, the receipt point delta is negative, the delivery point delta is positive, and both of them increase in their corresponding time t futures prices.

The structural properties in part (b) of this result formalize a natural behavior of the value of capacity in the receipt and delivery futures prices under model (2.10)-(2.14). The formulas in part (a) are obviously useful for computational purposes. Moreover, as shown below, they also offer an interesting operational and marketing interpretation of the delta positions $M_i(\cdot,T)$ and $N_j(\cdot,T)$.

Consider any time $t \in [0,T)$. Conditional on time t information, it is clear that random variables $\tilde{F}_i(T,T)/F_i(t,T)$ and $\tilde{G}_j(T,T)/F_i(t,T)$ have means equal to one under the risk neutral measure. Since these random variables are nonnegative, expressions (2.16)-(2.17) suggest that they can be interpreted as "random weights" for the random variables optimal amounts procured at receipt point i and marketed at delivery point j at time T, respectively, that is, $\sum_{j\in\mathcal{D}(i)} \tilde{x}_{ij}^*(T)$ and $\sum_{i\in\mathcal{R}(j)} \tilde{x}_{ij}^*(T)$. Thus, the quantities $\delta_i^F(t,T)$ and $\delta_j^G(t,T)$ can be interpreted as discounted expected weighted amounts of natural gas optimally procured and marketed at these points at time T. We make this interpretation precise below by introducing the concept of a *weighing* distribution.

Define $\mathbf{F}(\cdot,T) := (F_i(\cdot,T), i \in \mathcal{R})$ and $\mathbf{G}(\cdot,T) := (G_j(\cdot,T), j \in \mathcal{D})$, and recall that these are m and n dimensional vectors, respectively. Consider the $m+n$ dimensional random vector

$(\tilde{\mathbf{F}}(T,T), \tilde{\mathbf{G}}(T,T))$ and let vector u denote one of its realizations. Given $t \in [0,T)$, let $\psi_t(u)$ be the risk neutral joint probability density function (pdf) of this random vector under model (2.10)-(2.14) conditional on $(\mathbf{F}(t,T), \mathbf{G}(t,T))$. Let u_{-i} be vector u with the i-th of its first m components removed, and denote $\tilde{\mathbf{F}}_{-i}(T,T)$ random vector $\tilde{\mathbf{F}}(T,T)$ with its i-th component removed. Let $\psi_t(u_{-i}|u_i)$ be the risk neutral joint pdf of random vector $(\tilde{\mathbf{F}}_{-i}(T,T), \tilde{\mathbf{G}}(T,T))$ conditional on $F_i(T,T)$ *and* $(\mathbf{F}(t,T), \mathbf{G}(t,T))$, and $\psi_t^{F_i}(u_i)$ the risk neutral marginal pdf of random variable $\tilde{F}_i(T,T)$ given $F_i(t,T)$ under model (2.10)-(2.14). Define

$$w_t^{F_i}(u) := \frac{u_i \psi_t(u)}{F_i(t,T)}.$$

Since $w_t^{F_i}(u) \geqslant 0$ and $\tilde{F}_i(T,T)$ is a martingale under the risk neutral measure, it holds that

$$\begin{aligned}
\int_{\Re_+^{m+n}} w_t^{F_i}(u)du = \int_{\Re_+^{m+n}} \frac{u_i\psi_t(u)}{F_i(t,T)}du &= \int_{\Re_+} \frac{u_i}{F_i(t,T)} \left(\underbrace{\int_{\Re_+^{m+n-1}} \psi_t(u_{-i}|u_i)du_{-i}}_{1} \right) \psi_t^{F_i}(u_i)du_i \\
&= \frac{1}{F_i(t,T)} \underbrace{\int_{\Re_+} u_i \psi_t^{F_i}(u_i)du_i}_{F_i(t,T)} = 1.
\end{aligned}$$

Thus, we call $w_t^{F_i}(u)$ the $F_i(t,T)$-weighing joint pdf. Similarly we define $w_t^{G_j}(u) := u_i\psi_t(u)/G_j(t,T)$ and call it the $G_j(t,T)$-weighing joint pdf. Denote E_t^{w,F_i} and E_t^{w,G_j} expectations taken using $w_t^{F_i}(u)$ and $w_t^{G_j}(u)$, respectively. Then, for $\forall i \in \mathcal{R}$, we obtain

$$\begin{aligned}
\delta_i^F(t,T) &= -\frac{e^{-r(T-t)}}{F_i(t,T)(1-\phi)} E_t \left[\tilde{F}_i(T,T) \sum_{j\in\mathcal{D}(i)} \tilde{x}_{ij}^*(T) \right] \\
&= -e^{-r(T-t)} \int_{\Re_+^{m+n}} \left[\sum_{j\in\mathcal{D}(i)} \frac{\tilde{x}_{ij}^*(T)}{1-\phi_{ij}} \frac{u_i\psi_t(u)}{F_i(t,T)} \right] du \\
&= -e^{-r(T-t)} \int_{\Re_+^{m+n}} \left[\sum_{j\in\mathcal{D}(i)} \frac{\tilde{x}_{ij}^*(T)}{1-\phi_{ij}} w_t^{F_i}(u) \right] du \\
&= -e^{-r(T-t)} E_t^{w,F_i} \left[\sum_{j\in\mathcal{D}(i)} \frac{\tilde{x}_{ij}^*(T)}{1-\phi_{ij}} \right].
\end{aligned}$$

Applying a similar logic to $\delta_j^G(t,T)$ yields that $\delta_j^G(t,T) = e^{-r(T-t)} E_t^{w,G_j} \left[\sum_{i\in\mathcal{R}(j)} \tilde{x}_{ij}^*(T) \right]$. Thus, $M_i(t,T)$ and $N_j(t,T)$ are the discounted expected quantities of natural gas optimally procured at location i and marketed at location j, respectively, at time T, with expectations taken under the $F_i(t,T)$-weighing and $G_j(t,T)$-weighing pdfs.

Table 2.3: Definitions of the Greeks in our natural gas transport valuation problem

Name	Description
Delta	The rate of change in the contract value with an initial price
Vega	The rate of change in the contract value with a price volatility
Gamma	The rate of change of delta with an initial price
Eta	The rate of change in the contract value with a correlation between two prices
Theta	The rate of change of the contract value with the passage of time

The managerial implication of this analysis is that delta hedging defines dynamic procurement and marketing policies that are consistent with the optimal exercise of the network contract real option. In other words, if a shipper delta hedges and takes delivery of the long futures positions and delivers under the short futures positions at maturity time, then this player will receive at each receipt point exactly the natural gas procured there under an optimal exercise policy. Once shipped, this natural gas also matches that needed to satisfy the short futures positions at each delivery point, which are also equal to the amount of natural gas marketed under an optimal exercise policy. (These interpretations remain valid even if the shipper closes the relevant futures positions and procures and sells natural gas on the spot market.) These arguments put on solid theoretical grounds a common interpretation among natural gas merchants of the deltas of simpler spread option based models, that is, as "expected" amounts procured/marketed at the capacity usage time.

2.6 Computing the Other Greeks

An important advantage of using the pathwise method to compute the deltas is that we can calculate the deltas and the contract value at the same time. However, it is well known that the pathwise method can not be used to compute Gammas, which are the second derivatives of the contract value with respect to the initial prices. Thus, we use the central difference method to compute the other Greeks defined in Table 2.3 (Hull 2000 [30]). Of course, we can also use the central difference method to compute the deltas. According to Broadie and Detemple (2004) [6], the central difference method and the pathwise method, when the latter is applicable, provide very close derivative estimates.

Denote $V(t,T;\theta)$ the dependence of $V(t,T)$ on scalar parameter θ. As before, $V_\theta(t,T)$ denotes the first derivative of $V(t,T)$ with respect to θ, and we indicate by $V_{\theta,\theta}(t,T)$ the second derivative of $V(t,T)$ with respect to θ: $V_\theta(t,T) \equiv \frac{\partial V(t,T;\theta)}{\partial \theta}$, $V_{\theta,\theta}(t,T) \equiv \frac{\partial^2 V(t,T;\theta)}{\partial \theta^2}$. In §2.7,

we estimate these quantities for each relevant parameter, using the following expressions:

$$\hat{V}_{\theta}(t,T) = \frac{V(\theta+h)-V(\theta-h)}{2h}, \tag{2.18}$$

$$\hat{V}_{\theta,\theta}(t,T) = \frac{V(\theta+h)-2V(\theta)+V(\theta-h)}{4h^2} \tag{2.19}$$

where h is a "small" positive rea number. It is important to point out that central difference estimates are biased (Glasserman 2004 [25]).

2.7 Empirical Results

In this section we present our empirical results on the valuation and delta hedging performance of our LPS model relative to those of our L/UB models and to the perfect delta hedging case, respectively, and numerically compute and investigate the Greeks. Before discussing these results, we illustrate the data used in our numerical experiments and their setting.

Data description and price model estimation. We base our numerical experiments on the Transco pipeline system, which extends from Texas to New York (NY) City, New Jersey, and Pennsylvania. It has six pricing zones. Here we focus on zones 1-4. Zones 1 and 2 are in Texas, Zone 3 covers Louisiana and Mississippi, and Zone 4 is in Alabama. For these zones and Henry Hub, Louisiana, we obtained daily closing spot prices covering the period 1/2/2001-12/1/2006 and futures prices for December 2006 delivery during the period 11/28/2005-11/28/2006 (11/28/2006 was the closing date of the December 2006 NYMEX natural gas futures contract). The spot prices for zones 1-4 and Henry Hub correspond to the spot prices reported by Gas Daily and are available, for example, from Platts. We obtained futures prices for Henry Hub from NYMEX and for the other zones from an energy trading company.

Figures 2.3 and 2.4 display the spot and futures price data sets, respectively. The spot prices at different locations appear to move closely together. The futures prices exhibit more heterogeneity as the time to maturity increases. We assume that the natural logarithms of the spot prices at these locations follow correlated mean reverting processes (see §2.5). We estimate their parameters using a standard linear regression method based on the spot price data (Clewlow and Strickland 2000 [11], pp.28-29). In this case, recall from §2.5 that under the risk neutral measure the time t volatility function of the maturity T futures price at a given location i is $\sigma_i(t,T) = \sigma_i \exp[-\kappa_i(T-t)]$ where κ_i and σ_i, respectively, are the mean reversion rate and volatility of the natural logarithm of the spot price at location i (here and below the superscripts used in §2.5 on the quantities σ_i, κ_i, and ρ_{ij} are omitted).

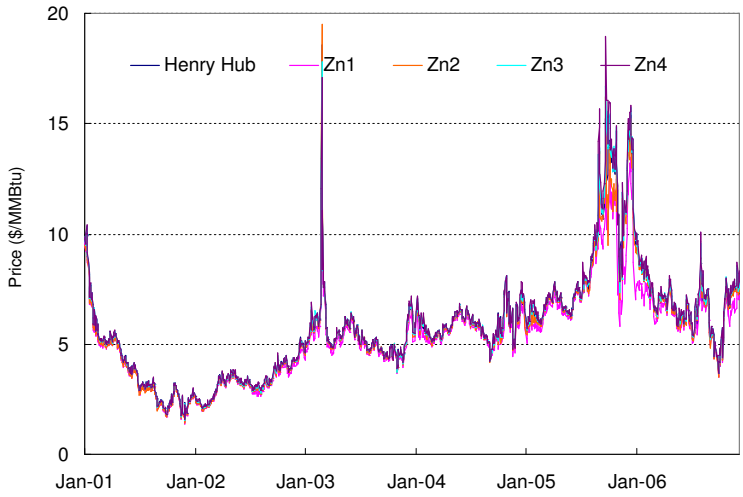

Figure 2.3: Relevant spot price data.

Moreover, the instantaneous correlation between the standard Brownian motions that drive the evolution of the two futures prices at locations i and j is ρ_{ij}, which is also the correlation between those associated with the natural logarithms of the spot prices at locations i and j. Since we have available futures prices at each location of interest and we work with the risk neutral measure, the only parameters we need for valuation and replication purposes are κ_i, σ_i, and ρ_{ij} with i, $j \in \{\text{Henry Hub, Zones 1-4}\}$ (if $i = j$ then $\rho_{ij} = 1$).

Table 2.4 reports the estimates of these parameters. A 95% confidence interval for the mean reversion rate for Henry Hub $\hat{\kappa}_H$ is $[1.864, 2.083]$. Similarly, the 95% confidence intervals for the mean reversion rates for Zones 1-4 are $[2.571, 2.819]$, $[2.639, 2.899]$, $[2.123, 2.357]$, and $[2.143, 2.377]$, respectively. Since we use linear regression to estimate the price parameters, we do not have the standard errors, and therefore the confidence intervals, for the relevant price volatilities and correlations.

We use the estimates, reported in Table 2.4, in the experiments described below. Henry Hub has smaller mean reversion rate and spot volatility estimates that the other markets. The estimates of the correlation coefficients are almost all above 90% and reflect the mentioned behaviors of the spot prices during the observation period.

Figure 2.5 displays the open interests and trading volumes for the futures markets of Henry Hub and Zones 3-4 (we do not have these data for Zones 1 and 2). Table 2.5 reports the trading volume statistics for each spot market. As expected, Henry Hub is the most liquid point both in terms of futures and spot markets. Zone 3 comes in second, while these figures are clearly lower in the other markets. However, a significant amount of trading also occurs on ICE and OTC markets, and Figure 2.5 does not include data associated

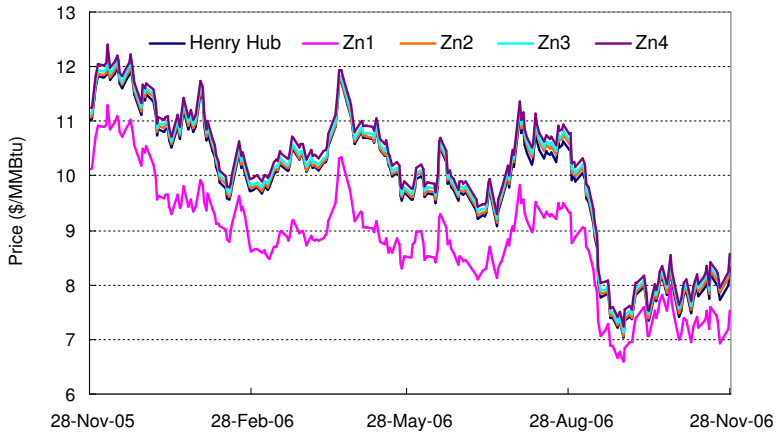

Figure 2.4: Relevant December 2006 maturity futures price data.

Table 2.4: Price model parameter estimates.

	Henry Hub	Zone 1	Zone 2	Zone 3	Zone 4
Mean reversion rate (κ)	1.974	2.695	2.769	2.240	2.260
Standard error of estimate of κ	0.056	0.063	0.066	0.060	0.060
Spot volatility	0.854	0.927	0.988	0.914	0.925
Correlation Matrix					
	Henry Hub	Zone 1	Zone 2	Zone 3	Zone 4
Henry Hub	1	0.906	0.932	0.955	0.948
Zone 1	0.906	1	0.893	0.910	0.912
Zone 2	0.932	0.893	1	0.925	0.927
Zone 3	0.955	0.910	0.925	1	0.982
Zone 4	0.948	0.912	0.927	0.982	1

Table 2.5: Spot market trading volumes in our data set.

Daily Volume (1,000 MMBtu/day)	Henry Hub	Zone 1	Zone 2	Zone 3	Zone 4
Average	2,077	73	103	732	162
Maximum	11,133	436	497	2,529	971
Minimum	148	1	11	6	2

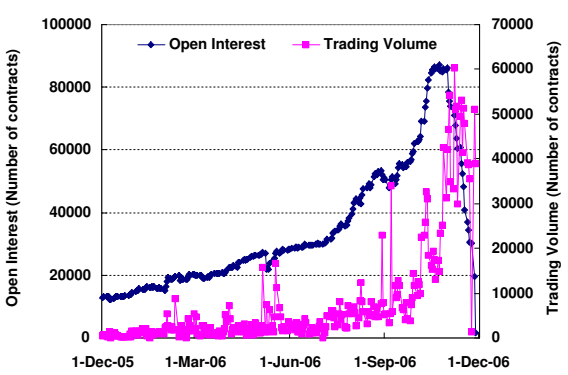

(a) Henry Hub

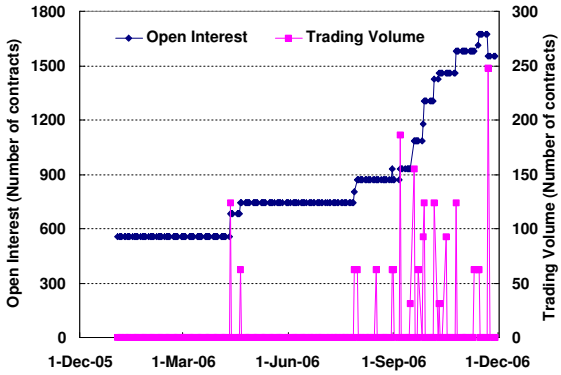

(b) Zone 3

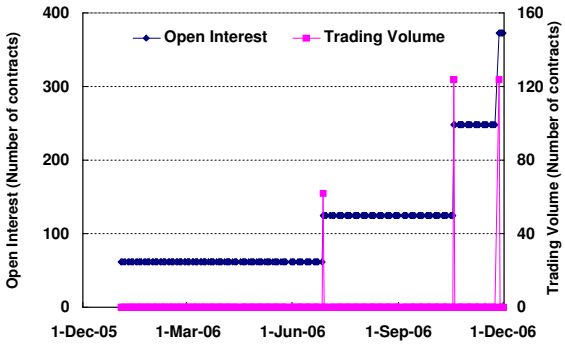

(c) Zone 4

Figure 2.5: Futures open interest and trading volume for the December 2006 maturity for select NYMEX markets (1 contract = 10,000 MMBtu).

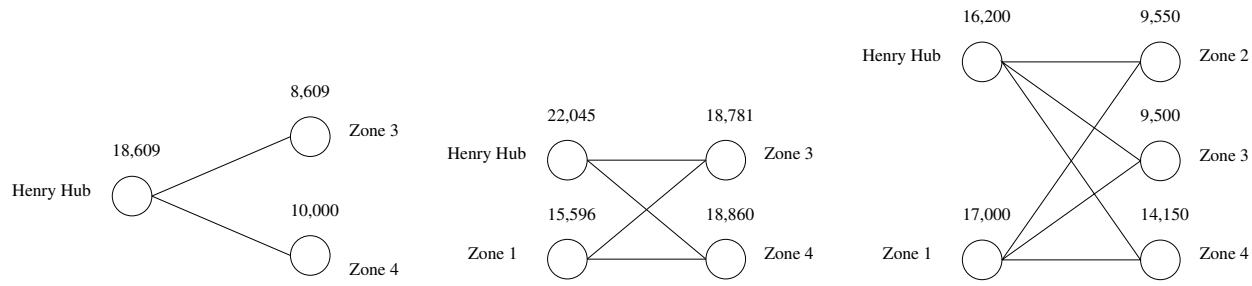

(a) 1R-2D (b) 2R-2D (c) 2R-3D

Figure 2.6: Contract networks used in our empirical study.

Table 2.6: Commodity and fuel rates (Source: Transco pipeline online information system).

Commodity Rate ($/MMBtu)					
	Henry Hub	Zone 1	Zone 2	Zone 3	Zone 4
Henry Hub	n.a.	0.00652	0.00507	0.00268	0.01372
Zone 1	0.00652	0.00162	0.00401	0.00652	0.01756
Zone 2	0.00507	0.00401	0.00251	0.00507	0.01611
Zone 3	0.00268	0.00652	0.00507	0.00268	0.01372
Zone 4	0.01372	0.01756	0.01611	0.01372	0.01121
Fuel Rate (%)					
	Henry Hub	Zone 1	Zone 2	Zone 3	Zone 4
Henry Hub	n.a.	1.05	0.78	0.39	2.14
Zone 1	1.05	0.27	0.66	1.05	2.80
Zone 2	0.78	0.66	0.39	0.78	2.53
Zone 3	0.39	1.05	0.78	0.39	2.14
Zone 4	2.14	2.80	2.53	2.14	1.75

with these trading activities. While different markets appear to exhibit different liquidity levels, in our experiments we assume that they are all equal. This is clearly a simplification, but accounting for heterogeneity in liquidity levels across markets would require estimating market specific transaction costs, such as bid-ask spreads, which goes beyond the scope of this paper.

Experiment setting. We consider the three contract networks shown in Figure 2.6, which we choose based upon examination of real contracts available on the Transco pipeline online information system. Network (a) models a contract with one receipt and two delivery points (1R-2D), (b) a contract with two receipt and two delivery points (2R-2D), and network (c) a contract with two receipt and three delivery points (2R-3D). The numbers associated

Table 2.7: Valuation results for the dedicated capacity case.

$T-t$	Contract Networks											
	1R-2D				2R-2D				2R-3D			
(month)	LB	UB	I	E	LB	UB	I	E	LB	UB	I	E
1	1	1	0.23	0.77	0.95	1.58	0.63	0.37	0.92	1.62	0.66	0.34
3	1	1	0.17	0.83	0.94	1.65	0.51	0.49	0.91	1.70	0.54	0.46
6	1	1	0.14	0.86	0.93	1.66	0.47	0.53	0.90	1.71	0.50	0.50
9	1	1	0.14	0.86	0.92	1.66	0.45	0.55	0.89	1.72	0.48	0.52
12	1	1	0.14	0.86	0.92	1.66	0.45	0.55	0.89	1.72	0.48	0.52

with each point are the point capacities, expressed in MMBtu/day. The term of each contract is the December 2006 month. We use the commodity and fuel rates reported in Table 2.6, which are available from the Transco pipeline online information system. We obtain the risk free rate from the relevant Treasury curve.

We solve the L/UB models by valuing each relevant spread options using Kirk's (1995 [37]) approximation formula (see Appendix 2.9.1). In the LB case, we also benchmarked our results when each relevant spread option is valued by simulation. Since we do not observe noticeable differences in the resulting valuations, we only report the results based on Kirk's formula. (We do not do this benchmarking in the UB case because we would have to simulate futures prices while simultaneously optimizing the spread option effective strike prices; in the LB case these strike prices are given before the optimization takes place.) We value each contract using our LPS model based on 100,000 samples for the futures prices in less that 1 Cpu second.

Valuation results. We first discuss the fully dedicated capacity case, then we consider three additional flexible capacity cases. We report the valuations yielded by our L/UB models as fractions of those of our LPS model. For the latter valuations we also display their relative intrinsic and extrinsic values. We do this for 1, 3, 6, 9, and 12 month(s) to contract expiration.

Dedicated capacity. Table 2.7 reports the relevant figures for this case. The L/UBs are tight in the 1R-2D contract cases. The result for the LB model is as expected, since a dedicated contract with one receipt point can be valued as a collection of spread options (see Table 2.2 in §2.1). In this simple case the UB is also tight. This network contract exhibits a large extrinsic value. This is due to the two spread options being almost at the money at contract inception, which implies that there is significant uncertainty as to whether these spreads would expire in or out of the money. As expected, the importance of the extrinsic

value increases with the number of months left until contract expiration.

The LB valuations of the 2R-2D and 2R-3D contracts are 5-8% and 8-11%, respectively, below those of the LPS model when the time to expiration is one month, and become somewhat looser when this time becomes longer. This is due to the fact that the LB model is unable to fully capture the extrinsic value embedded in these contracts. To illustrate, consider the 2R-2D contracts. The LB model assigns zero and positive flows, respectively, to the Zone 1-3 and Zone 1-4 spread options at each of the considered contract inception times. If the Zone 1-4 spread is out of the money at contract expiration, according to the LB model optimal solution no injection is made in Zone 1. However, it is possible that the Zone 1-3 spread is in the money at this time, and it may be optimal to ship natural gas from Zone 1 to Zone 3, which is feasible in our LPS model but impossible in the LB model. Moreover, even when both the Zone 1-3 and Zone 1-4 spread options expire in the money, it may be optimal to ship natural gas from Zone 1 to Zone 3, rather than from Zone 1 to Zone 4. Without fully capturing these features, the LB model underestimates the contract value.

The UB valuations of the 2R-2D and 2R-3D contracts are very loose, which shows that the UB models cannot be usefully employed for valuation purposes when a contract includes more than point in both the receipt and delivery sets.

Interestingly, the extrinsic values of the 2R-2D and 2R-3D contracts are lower than those in the 1R-2D contract case but are nevertheless significant, ranging from 34% to 55%. This is due to the fact that the Zone 1-4 spread is deep in the money at each contract inception time and will be used with high probability at the expiration time. Thus the contract intrinsic values play a larger role in cases 2R-2D and 2R-3D than in case 1R-2D.

Flexible capacity. We use the 2R-2D contracts to investigate how increasing capacity flexibility affects the comparisons between our L/UB and LPS models. We consider three additional cases: receipt flexibility, delivery flexibility, and receipt and delivery flexibility. In each of these cases, a shipper is free to use all the receipt/delivery capacity at a single point in its respective set of points. It is not clear that the relative performance of our LB model would deteriorate relative to our LPS model. With flexibility in the receipt or delivery capacity, at expiration it is optimal to procure or market natural gas from a single receipt or delivery point, and the LB model would do this. However, this model must select these points a priori and, consequently, may not yield the correct contract valuations. These aspects appear to affect the LB model performance in different directions and it does not seem clear which aspect would have relatively more importance. It is also not clear how different capacity flexibility structures would affect the contract intrinsic and extrinsic values.

Table 2.8 reports the relevant valuation results (this table includes the dedicated capacity case for ease of visual comparison). The LB model undervalues the contract by 6-13% with

Table 2.8: Valuation results under different capacity flexibility cases for the 2R-2D contract network.

Month(s) to Expiration	Dedicated				Receipt Flexibility			
	LB	UB	I	E	LB	UB	I	E
1	0.95	1.58	0.63	0.37	0.94	1.17	0.86	0.14
3	0.94	1.65	0.51	0.49	0.90	1.22	0.73	0.27
6	0.93	1.66	0.47	0.53	0.88	1.23	0.68	0.32
9	0.92	1.66	0.45	0.55	0.87	1.23	0.66	0.34
12	0.92	1.66	0.45	0.55	0.87	1.23	0.65	0.35
Month(s) to Expiration	Delivery Flexibility				Receipt and Delivery Flexibility			
	LB	UB	I	E	LB	UB	I	E
1	0.84	1.47	0.61	0.39	0.85	1.43	0.78	0.22
3	0.81	1.52	0.48	0.52	0.78	1.56	0.65	0.35
6	0.80	1.53	0.44	0.56	0.76	1.60	0.60	0.40
9	0.80	1.54	0.43	0.57	0.75	1.61	0.58	0.42
12	0.80	1.54	0.42	0.58	0.75	1.61	0.57	0.43

Table 2.9: LPS model delta hedging performance under different capacity flexibility cases for the 2R-2D contract network.

	Discrepancy (%)			
Month(s) to Expiration	Dedicated	Receipt Flexibility	Delivery Flexibility	Receipt and Delivery Flexibility
1	1.78	1.60	1.65	3.06
2	1.93	5.54	3.27	5.63
3	0.33	4.61	4.43	4.55
4	2.03	6.24	5.68	5.47
5	2.38	6.77	7.02	6.74
6	3.92	7.63	8.33	8.05
7	2.36	6.07	5.54	5.83
8	2.84	6.61	5.89	6.26
9	3.27	7.23	6.32	6.74
10	1.96	5.67	4.34	4.89
11	1.85	5.69	4.22	4.98
12	1.64	5.88	4.16	5.08

flexible receipt capacity, by 16-20% with flexible delivery capacity, and by 15-25% with flexible receipt and delivery capacity. Thus, the relative LB valuations decrease with capacity flexibility than in the dedicated capacity case, and these valuations are lower when flexibility is added to the delivery rather than the receipt capacity. This occurs because at contract inception the receipt point prices are more far apart from each other than are the delivery prices. For example, with 6 months to maturity, the Henry Hub and Zone 1 futures prices are \$9.758 and \$8.796, respectively, while the Zone 3 and Zone 4 futures prices are \$9.873 and \$9.963, respectively. Thus, the choice of which receipt point will be used at the contract expiration time with flexible receipt capacity seems relatively more clear than the choice of which delivery point will be used at this time with flexible delivery capacity. A similar argument also explains why the intrinsic value is relatively more important with flexible receipt capacity than in both the cases with dedicated and flexible delivery capacity.

Delta hedging. As explained in §2.5, delta hedging gives us a way to assess the accuracy of our LPS model by measuring what we named the discrepancy ratios. We focus on the 2R-2D contract under the four capacity flexibility cases previously discussed. We perform delta hedging on each of the Henry Hub and Zones 1, 3, and 4 futures markets once a day during the time period 12/1/2005-11/27/2006. To measure the effect of the time to expiration we compute twelve different contract values $V(0,T)$ each time by sliding the contract inception

date by one month, that is, we slide time 0 forward by one month each time. In each of these cases we also recompute the associated value for the expression $V(T,T) + \widehat{\text{CH}}(0,T)$.

Table 2.9 reports the discrepancy ratios $[V(0,T) - V(T,T) - \widehat{\text{CH}}(0,T)]/V(0,T)$ for each of the contract inception dates. These discrepancies range from 0.33% to 8.33%. These values do not seem to display a monotonic pattern as the time to expiration increases: if one were to superimpose a pattern on these values, they *roughly* increase up to 6 months to maturity and then decrease when the number of months to expiration increases further.

However, it does seem that the discrepancies are lower in the dedicated capacity case than in the flexible capacity cases. This can be explained as follows: as the futures prices change over time, based on the interpretation of the deltas established in §2.5, in the dedicated capacity case one would anticipate that the discounted expected weighted amounts of natural gas procured and marketed at the receipt and delivery points should be more stable than in the flexible capacity cases. Thus, one would expect the change in the contract value to be relatively more "predictable" and the discrepancies to be lower in the former case than the latter cases. Undr perfect delta hedging setting, the change of the contract value can be captured through delta hedging along the time. That is, the discrepancy of estimate value at time 0, $V(0,T)$ and total realized value at time T, $V(T,T) + \widehat{\text{CH}}(0,T)$, should be zero, if the contract value is perfectly hedged. The results in Table 2.9 indicate that the replicating performance of our LPS model is encouraging.

Greeks. We analyze the Greeks by only considering the 2R2D case. We let the time to maturity be 6 months. At this contract valuation time (t), the futures prices at Henry Hub, and Zones 1-4 are \$9.758, \$8.796, \$9.873, and \$9.963, respectively. We use the central difference method to estimate all the Greeks defined in Table 2.3. Table 2.10 displays the estimates of the Greeks. We also use the pathwise method to estimate the deltas. The deltas estimated using these two methods are very close, which is consistent with the results of Broadie and Detemple (2004) [6]. Therefore, using the central difference method and the pathwise method provides essentially equal delta hedging (replication) performance.

Table 2.10 shows that the contract value (a) is decreasing in the initial prices of the receipt points (negative deltas), and increasing in those of the delivery points (positive deltas); (b) is convex in the receipt and delivery points initial prices (positive Gammas); (c) is increasing in price volatilities (positive vegas); (d) is decreasing in the instantaneous price correlations (negative etas); and (e) for fixed contract maturity date, the contract value decreases when the time to maturity decreases (negative theta).

We further analyze the behavior of the contract value over a range of the relevant parameters in Figures 2.7-2.11 by varying one parameter at a time. That is, we conduct a comparative statics analysis. The vertical axis of these figures is per unit contract value

Greek	Central Difference Method		Pathwise Method	
	Estimate	Standard Deviation	Estimate	Standard Deviation
Henry Hub Delta	-0.27804	0.00035	-0.27802	0.00023
Zone 1 Delta	-0.26655	0.00039	-0.26655	0.00029
Zone 3 Delta	0.32965	0.00034	0.32966	0.00028
Zone 4 Delta	0.29502	0.00028	0.29503	0.00019
Henry Hub Gamma	0.46059	0.0024		
Zone 1 Gamma	0.13624	0.0019		
Zone 3 Gamma	0.04512	0.0012		
Zone 4 Gamma	0.30801	0.0023		
Henry Hub Vega	0.12855	0.00063		
Zone 1 Vega	0.01719	0.00064		
Zone 3 Vega	0.37144	0.00079		
Zone 4 Vega	0.43082	0.00081		
Henry Hub - Zone 3 Eta	-1.5103	0.0023		
Henry Hub - Zone 4 Eta	-1.5529	0.0032		
Zone 1 - Zone 3 Eta	-0.5540	0.0012		
Zone 1 - Zone 4 Eta	-0.4488	0.0010		
Henry Hub - Zone 1 Eta	-0.2432	0.0007		
Zone 3 - Zone 4 Eta	-1.1679	0.0021		
Theta	-0.1730	0.0009		

Table 2.10: Estimation of the Greeks

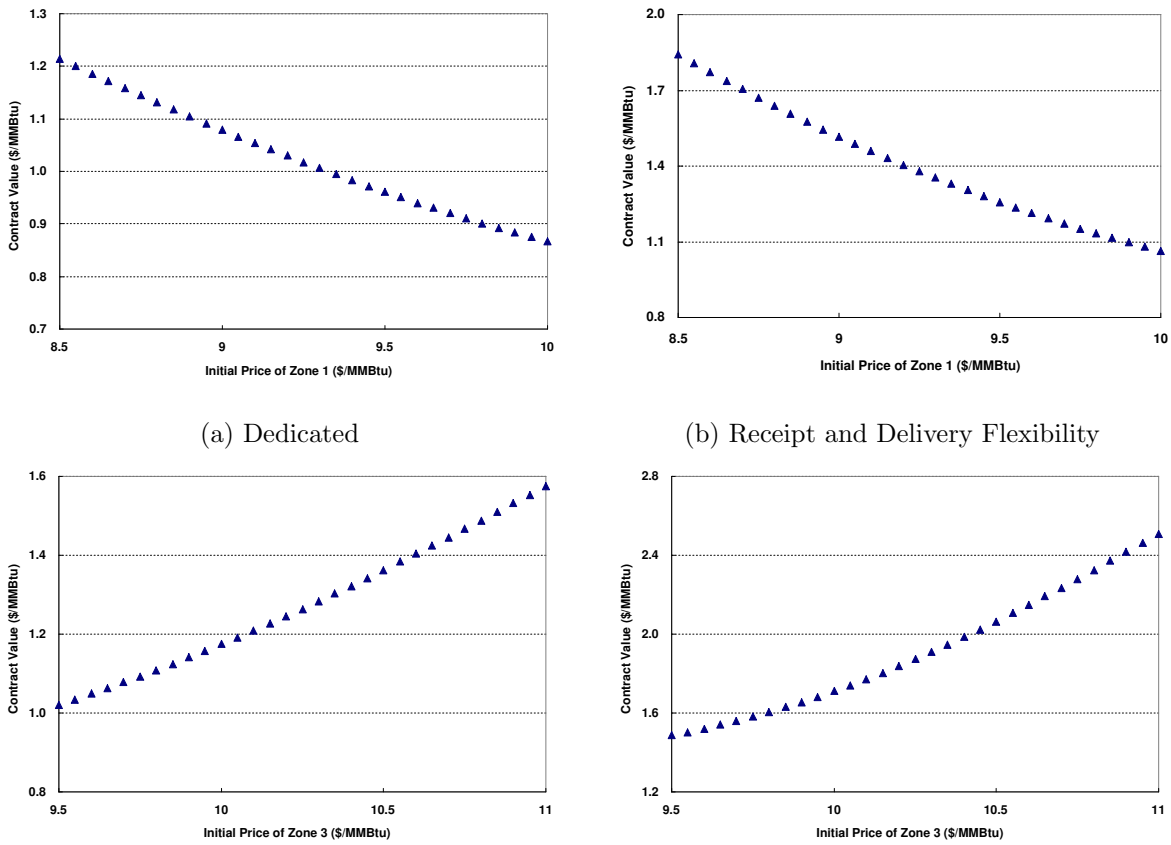

(a) Dedicated

(b) Receipt and Delivery Flexibility

(c) Dedicated

(d) Receipt and Delivery Flexibility

Figure 2.7: Comparative statics on initial prices

(unit: \$/MMBtu), that is, the contract value divided by the total contract capacity. We only consider the dedicated capacity and flexible receipt-and-delivery capacity cases. We obtain similar results in the remaining cases.

Initial Prices. We change the initial receipt price of Zone 1 or the initial delivery price of Zone 3. Figure 2.7 shows that the contract value is convex and decreasing in this initial receipt price, and convex and increasing in this delivery price, which is consistent with Proposition 4.

Price Volatilities. We change the price volatility of Zone 1 in two cases: (i) keeping the original correlations shown in Table 2.4, which are greater than 0.893; and (ii) setting all of these to zero. In case (i), panels (a) and (b) of Figure 2.8 show that the contract value is convex, decreasing first, and then increasing in this price volatility. In case (ii), panels (c) and (d) of Figure 2.8 show that the contract value is convex and increasing in this price volatility. Thus, how the contract value changes with volatility depends on the price correlations. We obtain similar results for the delivery price volatility.

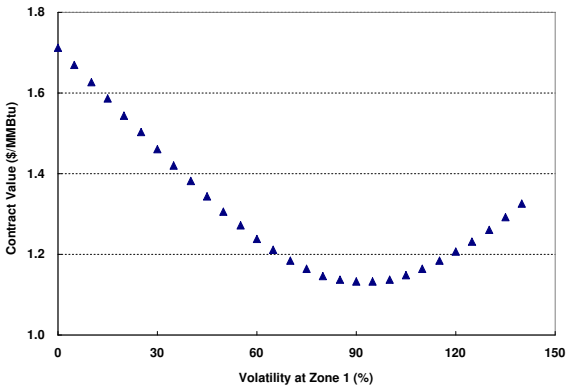

(a) Dedicated
(Correlation⩾0.893)

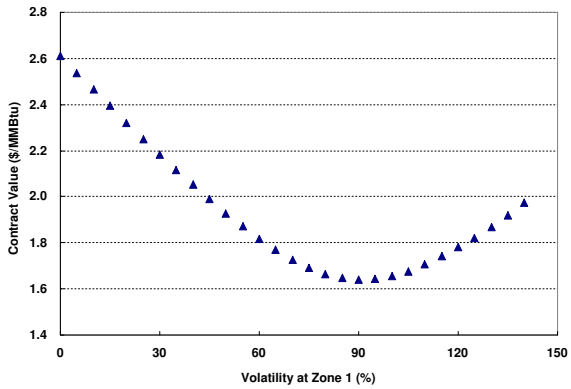

(b) Receipt and Delivery Flexibility
(Correlation⩾0.893)

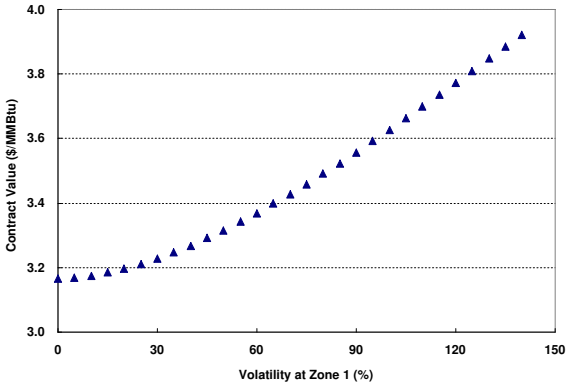

(c) Dedicated
(Correlation=0)

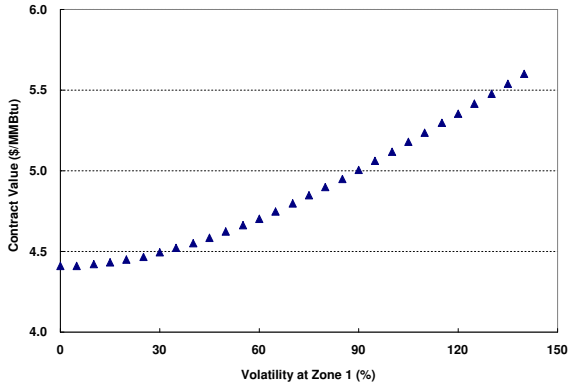

(d) Receipt and Delivery Flexibility
(Correlation=0)

Figure 2.8: Comparative statics on price volatilities

The "V" shape behavior of the contract values displayed in panels (a) and (b) of Figure 2.8 has also been observed by Eydeland and Wolyniec 2003 [21] for the point-to-point transport contract case. Here, we show that this effect persists in the network case.

Receipt-Delivery (Cross) Price Correlations. In this experiment we change one cross correlations and set the other correlations to zero. Panels Figure (a)-(d) and (e)-(h) of 2.9 display how the contract value changes in four cross correlations for the dedicated capacity and flexible receipt and delivery capacity cases, respectively. The results show that the contract values are decreasing in the cross correlations. This is intuitive, since the value of a spread option also decreases in the correlation between the prices of the underlying assets.

Receipt-Receipt ($\mathcal{R}$-$\mathcal{R}$) and Delivery-Delivery ($\mathcal{D}$-$\mathcal{D}$) Correlations. Figure 2.10 shows that the contract value also decreases in the $\mathcal{R}$-$\mathcal{R}$ and $\mathcal{D}$-$\mathcal{D}$ price correlations in both the dedicated and flexible receipt and delivery capacity cases.

Time to Maturity. Figure 2.11 shows that the contract value increases in a concave fashion, when the time to maturity increases, which is expected.

2.8 Conclusions

In this chapter we study the problem of computing and delta hedging the economic value of natural gas pipeline network transport capacity as a real option on natural gas prices at different geographical locations, and we use numerical examples to study how the contract value changes with the price parameters. The novelty of this problem resides in the network structure of capacity contracts, whereas the extant literature has only considered the point-to-point case. We develop and analyze an exact model based on linear programming and simulation that can be optimally solved very efficiently, and derive lower and upper bounds on the valuations yielded by this model. We study and interpret the financial hedging (delta) positions of this model as dynamic procurement/marketing policies that are consistent with the quantities of natural gas optimally purchased/sold at the contract execution time.

We use real data to test the valuation performance of our model against those of our bounding models, one of which is representative of current practice, and its delta hedging performance against the perfect replication benchmark, which is not achievable in practice. Our proposed model significantly outperforms current practice in terms of capacity valuations (by 5-25%), and exhibits encouraging delta hedging performance. Thus, this model emerges as a managerially relevant model for practical implementation and usage.

We also perform a comparative statics analysis of the contract value with respect to key price related parameters. The insights of this analysis are useful to support financial trading decisions.

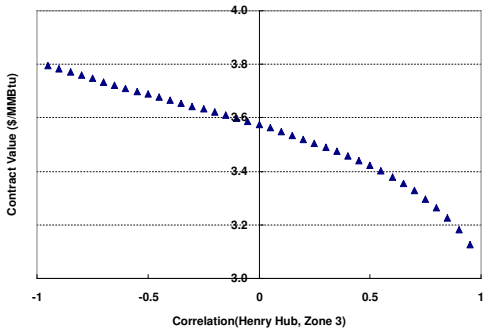

(a) Dedicated

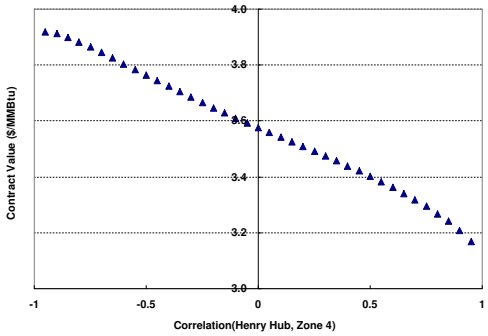

(b) Dedicated

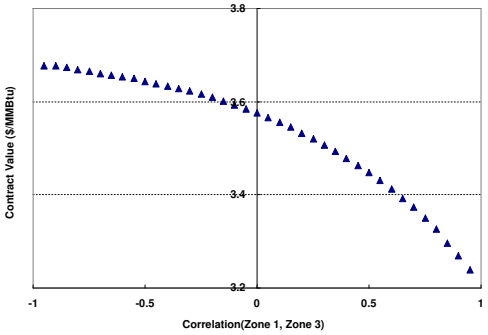

(c) Dedicated

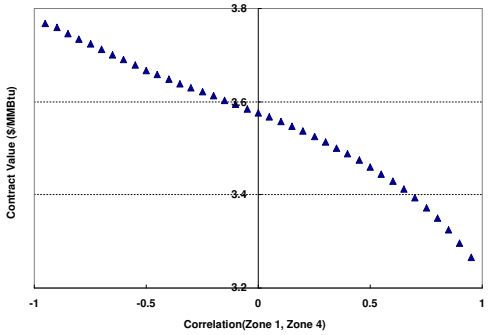

(d) Dedicated

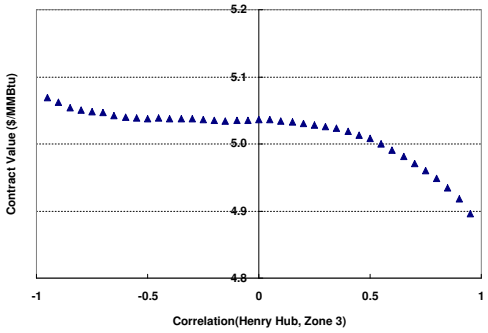

(e) Receipt and Delivery Flexibility

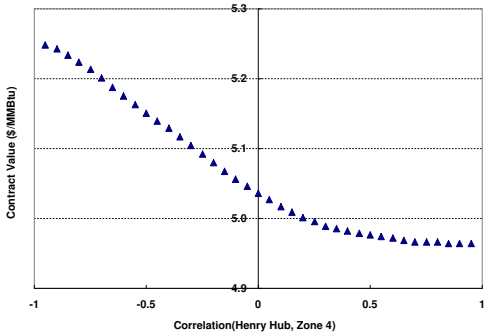

(f) Receipt and Delivery Flexibility

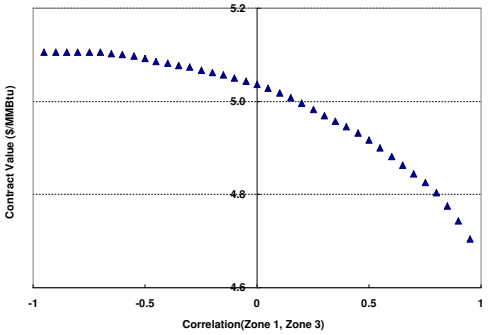

(g) Receipt and Delivery Flexibility

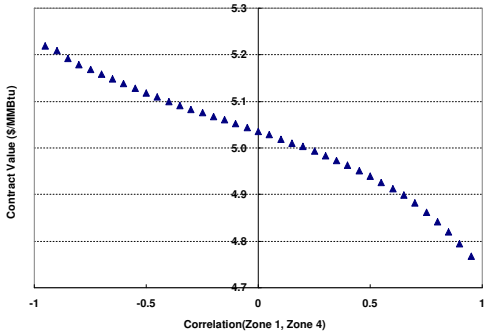

(h) Receipt and Delivery Flexibility

Figure 2.9: Comparative statics on cross price correlations

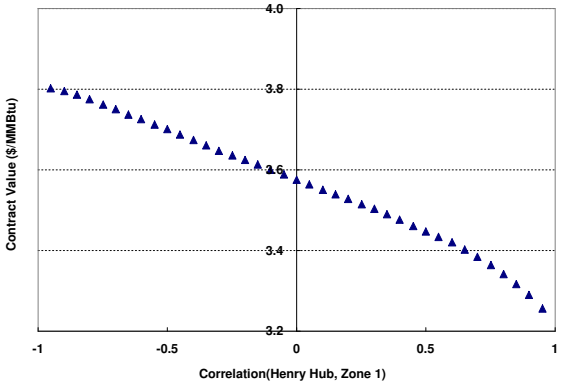

(a) Dedicated

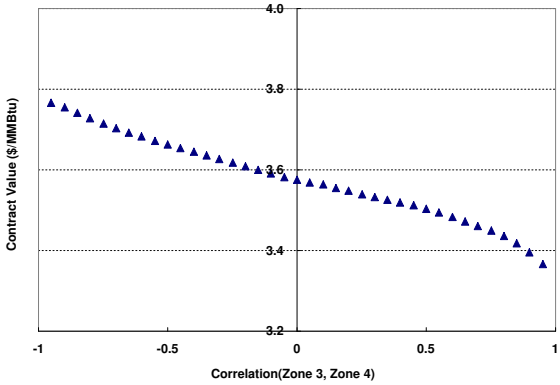

(b) Dedicated

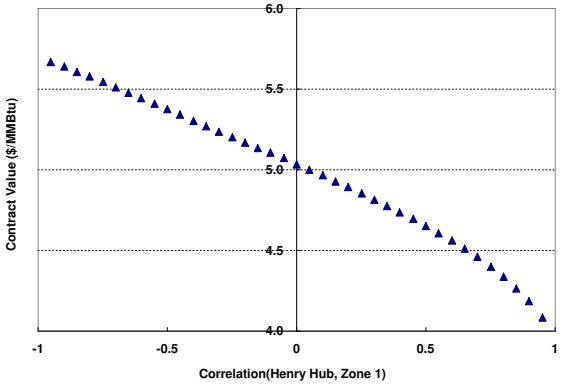

(c) Receipt and Delivery Flexibility

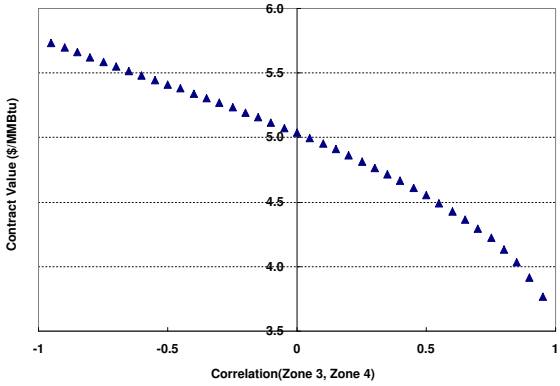

(d) Receipt and Delivery Flexibility

Figure 2.10: Comparative statics on $\mathcal{R}$-$\mathcal{R}$ and $\mathcal{D}$-$\mathcal{D}$ price correlations

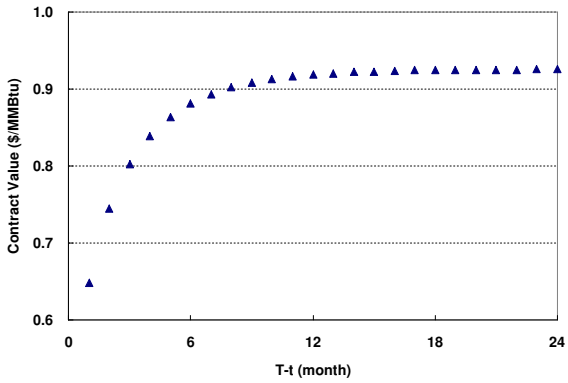

(a) Dedicated capacity

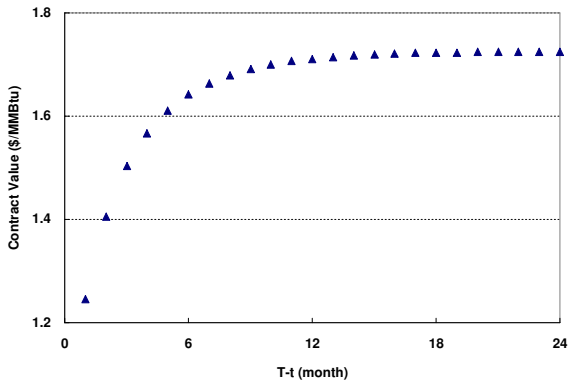

(b) Flexible receipt and delivery

Figure 2.11: Comparative statics on time to maturity

Our model can be extended to capture additional features of how natural gas pipeline capacity is sold via contracts in practice. These include right of first refusal clauses that give shippers the right to renew their transport contracts upon expiration before their capacity is sold to other shippers, and primary/secondary point and contingent capacity rights that give shippers different priorities and more flexibility, respectively, in the usage of contracted capacity. Moreover, our work can be further extended by empirically testing our model on different data sets, using alternative price models of the evolution of natural gas prices, and assessing the impact of bid/ask spreads on our model valuation/hedging performance; our neglecting of this impact is a limitation of our work.

Further applications beyond the natural gas industry are also of interest. Our work can be extended to deal with the real option valuation of refining, processing, shipping, and distribution capacity of other commodities, such as coal, oil, and biofuels. In these settings, modification of our model may include the addition of lead times.

2.9 Appendix of Chapter 2

2.9.1 Kirk's Spread Option Approximate Valuation Formula

Given receipt and delivery points i and j, suppose that random variables $\tilde{F}_i(T,T)$ and $\tilde{G}_j(T,T)$ are jointly lognormally distributed conditionally on their time 0 values $F_i(0,T)$ and $G_j(0,T)$ being known. In this case, there is no exact closed form expression for the time 0 value of payoff $\{\tilde{G}_j(T,T) - [F_i(0,T)/(1-\phi_{ij}) + K_{ij}]\}^+$ when $K_{ij} > 0$. Denote this value $V_{ij}^K(0,T)$. Kirk (1995 [37]) proposed the following approximation formula:

$$\begin{aligned}
V_{ij}^K(0,T) &= e^{-rT}\left\{G_j(0,T)\Phi(D_1) - \left[\frac{F_i(0,T)}{1-\phi_{ij}} + K_{ij}\right]\Phi(D_2)\right\} \\
s_{ij}^2 &:= \mathrm{Var}_0\left(\ln \tilde{G}_j(T,T) - \ln\left\{\tilde{F}_i(T,T)/[(1-\phi_{ij}) + K_{ij}]\right\}\right) \\
D_1 &:= \frac{\ln\{G_j(0,T)/[F_i(0,T)/(1-\phi_{ij}) + K_{ij}]\} + s_{ij}^2/2}{s_{ij}} \\
D_2 &:= D_1 - s_{ij},
\end{aligned}$$

where Var_0 is conditional variance given $F_i(0,T)$ and $G_j(0,T)$ and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal random variable. When $K_{ij} = 0$, this formula reduces to Margrabe's (1978 [41]) exact formula (see Carmona and Durrleman 2003 [9]).

2.9.2 Proof of Part (b) of Proposition 4

In futures price model (2.10)-(2.14), $\sigma_i^F(t,T)$, $\sigma_j^G(t,T)$, $\rho_{ii'}^{FF}(t)$, $\rho_{jj'}^{GG}(t)$, and $\rho_{ij}^{FG}(t)$ are not price dependent. This allows us to write

$$\begin{aligned}
\tilde{F}_i(T,T) &= F_i(0,T)\tilde{a}_i(\sigma_i^F,0,T) \\
\tilde{a}_i(\sigma_i^F,0,T) &:= \exp\left[-\frac{1}{2}\int_0^T(\sigma_i^F(u,T))^2du+\int_0^T\sigma_i^F(u,T)dW_i(u)\right] \\
\tilde{G}_j(T,T) &= G_j(0,T)\tilde{b}_j(\sigma_j^G,0,T) \\
\tilde{b}_j(\sigma_j^G,0,T) &:= \exp\left[-\frac{1}{2}\int_0^T(\sigma_j^G(u,T))^2du+\int_0^T\sigma_j^G(u,T)dW_j(u)\right].
\end{aligned}$$

Without loss of generality, focus on time 0. Consider the behavior of $V(0,T)$ in initial delivery price $G_j(0,T)$. To simplify the exposition, remove suffix $(0,T)$ from $G_j(0,T)$. Also remove one of the two T's from the arguments of $G_j(T,T)$ and $F_i(T,T)$. Notice that given G_j, it holds that $\tilde{G}_j(T)=G_j\tilde{b}_j(\sigma_j^G,T)$, where $\tilde{b}_j(\cdot)$ is a lognormal random variable that does not depend on G_j. Fix arbitrary $\overline{j}\in\mathcal{D}$ and pick $G_{\overline{j}}^2>G_{\overline{j}}^1$. Denote $(b_j(\cdot),\ j\in\mathcal{D})$ and $(F_i(T),\ i\in\mathcal{R})$ realizations of random vectors $(\tilde{b}_j(\cdot),\forall j\in\mathcal{D})$ and $(\tilde{F}_i(T,T),\forall i\in\mathcal{R})$. Denote $\mathbf{G}_{-\overline{j}}:=(G_j,\forall j\in\mathcal{D}\setminus\{\overline{j}\})$ so that $(G_{\overline{j}},\mathbf{G}_{-\overline{j}})\equiv(G_j,\forall j\in\mathcal{D})$. Indicate by $v(T;G_{\overline{j}},\mathbf{G}_{-\overline{j}})$ and $x_{ij}^*(G_{\overline{j}},\mathbf{G}_{-\overline{j}})$ the dependence of $v(T)$ and x_{ij}^* on $(G_{\overline{j}},\mathbf{G}_{-\overline{j}})$, where x_{ij}^* is an optimal solution to model P. We have

$$\begin{aligned}
v(T;G_{\overline{j}}^2,\mathbf{G}_{-\overline{j}}) &= \sum_{i\in\mathcal{R}(\overline{j})}[G_{\overline{j}}^2b_{\overline{j}}(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^2,\mathbf{G}_{-\overline{j}}) \\
&\quad+\sum_{j\in\mathcal{D}\setminus\{\overline{j}\}}\sum_{i\in\mathcal{R}(j)}[G_jb_j(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^2,\mathbf{G}_{-\overline{j}}) \\
&\geqslant \sum_{i\in\mathcal{R}(\overline{j})}[G_{\overline{j}}^2b_{\overline{j}}(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}) \\
&\quad+\sum_{j\in\mathcal{D}\setminus\{\overline{j}\}}\sum_{i\in\mathcal{R}(j)}[G_jb_j(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}) \\
&\quad;\ \text{by optimality of } x_{ij}^*(G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}) \\
&\geqslant \sum_{i\in\mathcal{R}(\overline{j})}[G_{\overline{j}}^1b_{\overline{j}}(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}) \\
&\quad;\ \text{since } G_{\overline{j}}^2>G_{\overline{j}}^1 \\
&\quad+\sum_{j\in\mathcal{D}\setminus\{\overline{j}\}}\sum_{i\in\mathcal{R}(j)}[G_jb_j(\cdot)-\frac{F_i(T)}{1-\phi_{ij}}-K_{ij}]x_{ij}^*(G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}) \\
&= v(T;G_{\overline{j}}^1,\mathbf{G}_{-\overline{j}}).
\end{aligned}$$

Since this inequality holds for all realizations $(b_j(\cdot), \forall j \in \mathcal{D})$ and $(F_i(T), \forall i \in \mathcal{R})$, it follows that

$$V(0,T;G^2_{\bar{j}},\mathbf{G}_{-\bar{j}}) = e^{-rT}E_0[\tilde{v}(T;G^2_{\bar{j}},\mathbf{G}_{-\bar{j}})] \geqslant e^{-rT}E_0[\tilde{v}(T;G^1_{\bar{j}},\mathbf{G}_{-\bar{j}})] = V(0,T;G^1_{\bar{j}},\mathbf{G}_{-\bar{j}}),$$

and $V(0,T)$ increases in $G_{\bar{j}}$. Pick $\varphi \in (0,1)$ and $\bar{j} \in \mathcal{D}$ and define $G^{\varphi}_{\bar{j}} = \varphi G^1_{\bar{j}} + (1-\varphi)G^2_{\bar{j}}$. Notice that

$$\begin{aligned}
v(T;G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}}) &= \sum_{i\in\mathcal{R}(\bar{j})}[G^{\varphi}_{\bar{j}}b_{\bar{j}}(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}}) \\
&\quad + \sum_{j\in\mathcal{D}(i)\setminus\{\bar{j}\}}\sum_{i\in\mathcal{R}}[G_jb_j(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}}) \\
&= \varphi\left\{\sum_{i\in\mathcal{R}(\bar{j})}[G^1_{\bar{j}}b_{\bar{j}}(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}})\right. \\
&\qquad \left. + \sum_{j\in\mathcal{D}(i)\setminus\{\bar{j}\}}\sum_{i\in\mathcal{R}}[G_jb_j(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}})\right\} \\
&\quad +(1-\varphi)\left\{\sum_{i\in\mathcal{R}(\bar{j})}[G^2_{\bar{j}}b_{\bar{j}}(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}})\right. \\
&\qquad \left. + \sum_{j\in\mathcal{D}(i)\setminus\{\bar{j}\}}\sum_{i\in\mathcal{R}}[G_jb_j(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^{\varphi}_{\bar{j}},\mathbf{G}_{-\bar{j}})\right\} \\
&\leqslant \varphi\left\{\sum_{i\in\mathcal{R}(\bar{j})}[G^1_{\bar{j}}b_{\bar{j}}(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^1_{\bar{j}},\mathbf{G}_{-\bar{j}})\right. \\
&\qquad \left. + \sum_{j\in\mathcal{D}(i)\setminus\{\bar{j}\}}\sum_{i\in\mathcal{R}}[G_jb_j(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^1_{\bar{j}},\mathbf{G}_{-\bar{j}})\right\} \\
&\quad +(1-\varphi)\left\{\sum_{i\in\mathcal{R}(\bar{j})}[G^2_{\bar{j}}b_{\bar{j}}(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^2_{\bar{j}},\mathbf{G}_{-\bar{j}})\right. \\
&\qquad \left. + \sum_{j\in\mathcal{D}(i)\setminus\{\bar{j}\}}\sum_{i\in\mathcal{R}}[G_jb_j(\cdot) - \frac{F_i(T)}{1-\phi_{ij}} - K_{ij}]x^*_{ij}(G^2_{\bar{j}},\mathbf{G}_{-\bar{j}})\right\} \\
&= \varphi v(T;G^1_{\bar{j}},\mathbf{G}_{-\bar{j}}) + (1-\varphi)v(T;G^2_{\bar{j}},\mathbf{G}_{-\bar{j}})
\end{aligned}$$

This then implies that

$$\begin{aligned}
V(0,T;\mathbf{G}^{\varphi}) = e^{-rT}E_0[\tilde{v}(T;\mathbf{G}^{\varphi})] &\leqslant e^{-rT}\varphi E_0\tilde{v}(T;G^1_{\bar{j}},\mathbf{G}_{-\bar{j}}) + (1-\varphi)E_0\tilde{v}(T;G^2_{\bar{j}},\mathbf{G}_{-\bar{j}}) \\
&= \varphi V(0,T;G^1_{\bar{j}},\mathbf{G}_{-\bar{j}}) + (1-\varphi)V(0,T;G^2_{\bar{j}},\mathbf{G}_{-\bar{j}}),
\end{aligned}$$

and $V(0,T)$ is convex in $G_{\bar{j}}$.

Consider the behavior of $V(0,T)$ in $F_i(0,T)$. Make simplifications to the notation relative to $F_i(0,T)$ and $F_i(T,T)$ analogous to those made to $G_j(0,T)$ and $G_j(T,T)$. Fix arbitrary $\bar{i} \in \mathcal{R}$ and pick $F^2_{\bar{i}} > F^1_{\bar{i}}$. It holds that

$$\begin{aligned}
v(T; F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) &= \sum_{j \in \mathcal{D}(\bar{i})} [G_j(T) - \frac{F^2_{\bar{i}} a_{\bar{i}}(\cdot)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \\
&\quad + \sum_{i \in \mathcal{R} \setminus \{\bar{i}\}} \sum_{j \in \mathcal{D}(i)} [G_j(T) - \frac{F_i a_i(\cdot)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \\
&\leqslant \sum_{j \in \mathcal{D}(\bar{i})} [G_j(T) - \frac{F^1_{\bar{i}} a_{\bar{i}}(\cdot)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \\
&\quad + \sum_{i \in \mathcal{R} \setminus \{\bar{i}\}} \sum_{j \in \mathcal{D}(i)} [G_j(T) - \frac{F_i a_i(\cdot)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \\
&\leqslant \sum_{j \in \mathcal{D}(\bar{i})} [G_j(T) - \frac{F^1_{\bar{i}} a_{\bar{i}}(\cdot)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^1_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \\
&\quad + \sum_{i \in \mathcal{R} \setminus \{\bar{i}\}} \sum_{j \in \mathcal{D}(i)} [G_j(T) - \frac{F_i a_i(\epsilon_i)}{1-\phi_{ij}} - K_{ij}] x^*_{ij}(F^1_{\bar{j}}, \mathbf{F}_{-\bar{j}}),
\end{aligned}$$

which then implies $V(0,T; F^2_{\bar{j}}, \mathbf{F}_{-\bar{j}}) \leqslant V(0,T; F^1_{\bar{j}}, \mathbf{F}_{-\bar{j}})$, so that $V(0,T)$ decreases in F_j. The proof of the convexity of $V(0,T)$ in F_i is analogous to that of the convexity of $V(0,T)$ in G_j. □

Chapter 3

Valuation of Downstream Liquefied-Natural-Gas Storage

3.1 Introduction

Liquefied natural gas (LNG) is natural gas cooled to liquid state at -260F; liquefaction reduces the volume of natural gas by a factor of more than 600, making storage and shipping practical. Special ocean-going vessels load LNG at liquefaction facilities (for example in Algeria, Trinidad and Tobago, Australia or Qatar), transport it (for days or weeks), and unload it at terminals (for example in the U.S., Spain or Japan). At these terminals LNG is pumped into storage tanks, regasified, and then distributed via pipelines or, sometimes, by trucks. The Energy Information Administration (EIA) forecasts that local production will soon be unable to meet demand in most industrialized countries, and expects LNG imports to play an important role in bridging this gap (EIA 2003 [18], 2006 [20]). This increase in the world's natural-gas demand is mainly due to the fact that natural gas is an environmentally clean and abundant fuel, which has helped to make it the fuel of choice for many new power generation projects. More specifically, EIA (2006 [20]) forecasts that by 2010 LNG will become the largest source of U.S. natural-gas net imports.

We are starting to see the unfolding of some of these predicted events, which Jensen (2003 [32]) emphatically refers to as the "LNG revolution". Several liquefaction capacity expansion and greenfield projects have been announced, and a number of new terminals have been proposed in North America (see the website of the Federal Energy Regulatory Commission, www.ferc.gov/industries/lng.asp, for an up-to-date list of existing, approved, and proposed terminals in North America). An important feature of these new terminals is their larger sizes relative to the existing terminals, which should increase the ability of industry players to store LNG before releasing it into the natural-gas distribution system. Holcomb (2006 [29])

has recently questioned the effectiveness of storing LNG at a downstream terminal rather than storing regasified LNG in natural-gas storage facilities. However, obtaining access to these terminals is necessary to bring LNG into the natural-gas distribution system, at least when using conventional regasification technology, and requires leasing their storage space and regasification capacity from their operating companies. Hence, industry players face the challenge of assessing the value of downstream-terminal leasing contracts, which amounts to valuing the real option to store LNG at a terminal. Given the discussion in Holcomb (2006 [29]), one can surmise that the value of this real option is not well-understood in practice (let alone theory).

Focus and research questions. Our interest in this paper is on the modeling and analysis of the value of LNG downstream storage, focusing on regasification facilities located in the U.S., where LNG is priced off the New York Mercantile Exchange (NYMEX) natural-gas futures and basis-swap contracts. Specifically, our objective is to answer the following research questions. (1) How can one model the operations of an LNG value chain, in particular those pertaining to downstream storage? (2) What is the structure of the optimal LNG inventory-management policy at the downstream facility, and how does it interact with the logistics of inbound shipping? (3) What is the value of the real option to store LNG at this facility, or, equivalently, what is the value of optimally managing the LNG inventory at this facility? (4) What is the significance of this value for different players in an LNG value chain, i.e., LNG integrated-producers and merchants?

Contributions and relevance. We contribute to the literature in the following ways. On the modeling side, we address question (1) in §3.2 by developing a tractable real-option model for the valuation of the operations of an LNG chain, with an emphasis on the valuation of downstream storage. We abstract away from the details of upstream liquefaction and shipping and represent *both* these activities as a closed queueing network (CQN) of ships, the *shipping model*, leveraging and extending the work of Koenigsberg and Lam (1976 [38]) by allowing for Coxian shipping times. This model reflects current practices in the LNG industry aimed at achieving high utilization of installed liquefaction and shipping capacity. In contrast to the existing literature, we develop an analytical approach, the *rolling-forward* method, to probabilistically incorporate the output of the CQN in a finite-horizon stochastic dynamic-program, the *inventory-release model*, which optimizes the inventory-management policy at the downstream location. Here, we model the evolution of the wholesale spot price of natural gas as an exogenous stochastic Markov process. As a whole, our model is extremely practical: the shipping model captures the stochastic nature of the shipping process (including congestion), and *decomposes* it from the inventory-release model by using the rolling-forward method to determine the probability distribution of the number of

unloaded cargos during an inventory-review period. This decomposition makes our model computationally tractable.

On the analysis side, we answer question (2) in §3.3 by establishing optimality of a basestock target policy that depends on the realization, at the inventory-review time, of the random state variable that models the evolution of the natural-gas spot price. This is fundamentally a result of the interplay between shipping logistics and finite terminal-space/regasification-capacity, reflected in a "kinked" inventory and released-quantity feasible region in the inventory-release model. The basestock targets are nontrivial, in the sense that when it is optimal to sell, it is not necessarily optimal to sell the entire available inventory. We decompose the amount of inventory optimally released in each period into a forced sale and an optional sale: the former addresses the need to make space for incoming cargos, the latter prevents costly future forced sales or takes advantage of a high natural-gas price. More specifically, the basestock target is a sell-down-to level for an optional sale *after* a forced sale, if any, has been performed. This is a target because it may not be reachable from some inventory level due to limited regasification capacity. These structural results clarify an apparent misconception among some practitioners on how to manage LNG storage, which is summarized by Holcomb (2006 [29]) when he states that "the liquid storage must be vaporized and depleted before each new cargo of LNG arrives." In light of our analysis, a more accurate statement is that *some*, but not necessarily all, of the stored LNG must be released before a new cargo arrives. In fact, deciding the optimal amount of LNG to regasify and sell is, in general, nontrivial.

We use our model to answer questions (3) and (4) in §3.4 by means of a numerical study. We consider a realistic LNG chain consisting of liquefaction in North Africa, shipping to Lake Charles, Louisiana, and regasification and sale into the Louisiana natural-gas wholesale market at the Henry-Hub spot price, where Henry Hub is the delivery location of NYMEX natural gas futures. We model this price using the one-factor mean-reverting model with seasonality of Jaillet et al. (2004 [31]), calibrated with actual prices of traded NYMEX natural-gas futures and options on futures. We use this model in conjunction with our valuation model to measure the value of downstream storage over a twenty-year horizon, its sensitivity to parameters of interest, and its relative benefit for different parties. More specifically, we consider three players: An integrated producer that operates the whole LNG chain (an unusual case in practice, but one which some producers are currently considering), and two types of LNG merchants, type one and two for short. The type-one merchant purchases LNG from a (nonintegrated) producer at its liquefaction facility and ships LNG to a downstream terminal, where it also manages LNG inventory. This merchant operates under a free on board (FOB) contract, whereby the change of LNG ownership occurs at

the upstream facility, and also operates a fleet of vessels to ship LNG to the downstream facility. The type-two merchant purchases LNG from a (nonintegrated) producer at the downstream facility. This merchant operates under an ex-ship contract, whereby ownership changes hands at the regasification terminal, and the producer operates the shipping part of the chain. We obtain the value of storage by comparing the valuations generated by our model with optimization of the inventory-release policy and without storage, respectively. While this value does not depend on the type of player under consideration, we show that the *relative* benefit of storage does.

We estimate the value of the real option to store during this twenty-year period to be in the range of approximately \$100-525M, depending on the available storage space. For our case study this value is \$104M for a system with a throughput of 0.74 million tons per annum (MTPA), a storage space of 3BCF, and a regasification capacity of 1BCF/day, and it is \$526M for a 3.68MTPA throughput, 15BCF storage space, and 1BCF/day regasification capacity system. Approximately 20-25% of this value is extrinsic, i.e., can be attributed to natural-gas price volatility (stochastic variability), while the remaining 75-80% is intrinsic, i.e, can be attributed to natural-gas price seasonality (deterministic variability). We quantify the relative benefit of storage to be small for the integrated producer and significant for the two merchants. We estimate that an integrated producer of a typical size LNG chain can increase the value of its operations by roughly 0.5-3% by optimally managing the inventory at the downstream terminal, whereas this benefit rises to about 4-15% for the two types of merchants (both larger and smaller values are possible for other configurations). Hence, while we agree with Holcomb (2006 [29]) that an LNG downstream terminal "is a delivery mechanism" and not a substitute for natural gas (as opposed to LNG) storage, we also show that optimal management of stored LNG, given the available storage space and regasification capacity, can significantly enhance the economic value of this terminal. Put another way, while optimal inventory-management cannot increase process capacity, proper management of LNG storage can add significant value to an LNG chain by exploiting the marked fluctuations in the price of natural gas.

We conclude in §3.5 by further discussing the relevance of our results, pointing out some limitations and possible extensions/applications of our model, and delineating further research avenues.

Related literature. Our work is related to the real-option literature dealing with applications in commodity and energy industries. Representative works in this area are the valuation papers by Smith and McCardle (1998 [57], 1999 [57]) and references therein, Jaillet et al. (2004 [31]), and Manoliu (2004 [40]), and the recent books by Clewlow and Strickland (2000 [11]), Eydeland and Wolyniec (2003 [21]), and Geman (2005 [24]). Schwartz (1997

[50]), Seppi (2002 [54]), Clewlow and Strickland (2000 [11], Chapters 6-8), and Eydeland and Wolyniec (2003 [21], Chapters 4 and 5) provide reviews on modeling energy prices as stochastic processes. To the best of our knowledge, this literature has not yet studied the valuation of LNG operations as we do here, e.g., Geman (2005 [24], pp. 246-249) only briefly describes the LNG problem setting. While scant, the literature features two important papers that address the modeling of the shipping stage of an LNG chain: Koenigsberg and Lam (1976 [38]) and Kaplan et al. (1972 [35]). These works, the former based on CQN methods, the latter on simulation, endeavor to answer the following purely operational questions: given a configuration of an LNG system, i.e., the fleet size and its composition, the operational features of the loading and unloading facilities, and the distance between these facilities, what is the throughput of the system, i.e., the average amount of LNG delivered from the liquefaction to the regasification facility in steady state, and what are the utilization rates of the loading and unloading facilities? These are important questions, which however ignore the broader issue of valuing, in monetary terms, the LNG flow from the liquefaction to the regasification facility via the shipping link. While one could obtain a rough estimate of this value by multiplying the system throughput by an "appropriate average" natural-gas price, this approach would nevertheless ignore the value of downstream storage. Instead, differently from the extant literature, we address the downstream LNG storage valuation problem by modeling the interplay between shipping logistics and inventory management, in the face of the uncertain evolution of the natural-gas spot price.

Our inventory-release model is related to the literature dealing with the control of the water level in dams. Hasofer (1966a, b [27] [28]) considers the infinite-depth case, Faddy (1974 [22]), Pliska (1975 [47]), and Cohen and Rubinovitch (1977 [12]) the finite-depth case, and Moran (1959 [45]) studies both cases. Our setting is different because our dynamic program has an additional state variable, tracking the evolution of the natural-gas spot price, which evolves as a Markov process. In this regard, our work is related to the recent paper by Secomandi (2007 [53]), who studies the problem of valuing commodity storage assets, which give their owners the ability to control both the rates of injection and withdrawal of the commodity into and out of the storage facility. Our model is different from his because our inflow rate is uncontrollable (random). This feature, combined with the finiteness of the buffer size, creates a kink in the inventory-release feasible region that differentiates our setting from his. Our results reflect this basic difference since, in contrast to this author, we establish existence of a price-independent sell-down-to-level for forced sales and a price-dependent sell-down-to-target for optional sales. Moreover, differently from Secomandi (2007 [53]), we conduct a model-based analysis of the value of LNG storage using realistic data.

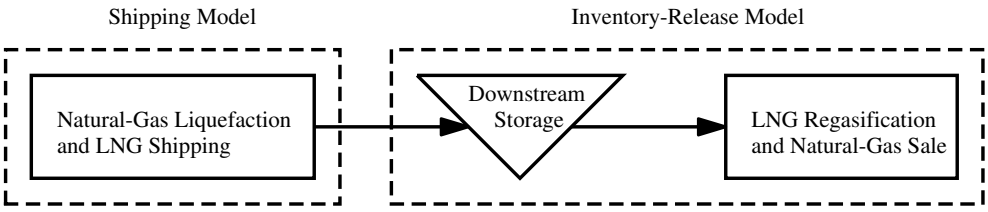

Figure 3.1: LNG process flow and modeling decomposition: The shipping model captures natural-gas liquefaction and LNG shipping, the inventory-release model management of downstream storage, LNG regasification, and natural-gas sale.

3.2 Model

In this section we describe our model for the valuation of LNG downstream storage, hence addressing research question (1). This model is based on the process-flow diagram of an LNG value chain illustrated in Figure 3.1. It consists of two integrated models: the shipping model (§3.2.1), a CQN-based model, and the inventory-release model (§3.2.2), a discrete-time finite-horizon stochastic dynamic-program. In Figure 3.1, there are two activities and one buffer. The first activity combines natural-gas liquefaction and LNG shipping; this is the shipping model. The buffer represents storage at the downstream regasification facility. The second activity includes LNG regasification and natural-gas sale into the wholesale market; this is the inventory-release model. The inventory-release model integrates the probabilistic output of the shipping model, generated by the rolling-forward method, with the management of the buffer and the second activity. It also uses the probabilistic output of a separate model of the evolution of the natural-gas spot price: a one-factor mean-reverting model. To maintain generality, we defer to §3.4 the discussion of the specifics of this model (the analysis in §3.3 also applies to other price models). In our model, we assume that the natural gas spot price and LNG import amount are independent. We obtain the value of storage by comparing the value of the LNG chain with optimization of the inventory-release policy to that without storage; we measure the relative benefit of storage for different players as the additional value created by this optimal policy relative to the value of their respective operations without storage (§3.2.3).

3.2.1 Shipping Model

CQN model. Following Koenigsberg and Lam (1976 [38]), we consider the CQN representation of the LNG shipping system illustrated in Figure 3.2. Ships perform four activities:

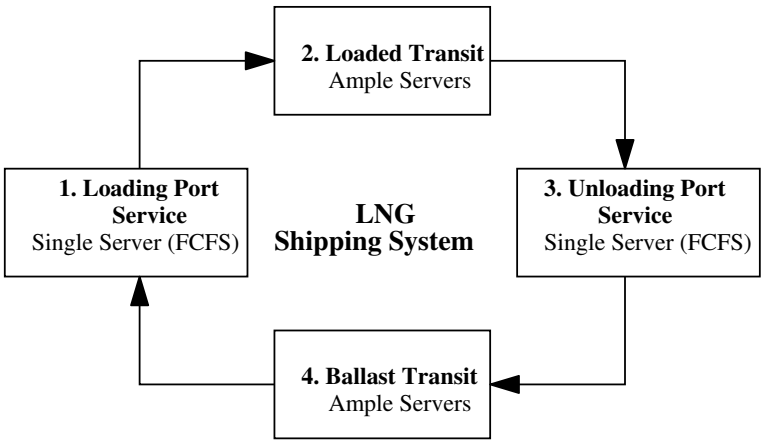

Figure 3.2: Representation of an LNG shipping system as a CQN of ships.

loading at the upstream port, loaded transit to the downstream port, unloading at this port, and ballast transit to the upstream port. We abstract away from the details of natural-gas liquefaction, and model the amount of LNG transfered from the liquefaction to the regasification facility by a fleet of N identical LNG ships, each with cargo size C measured in cubic meters (CM) fed by an ample supply. As LNG liquefaction facilities are designed to run at full capacity, being served by an appropriate number of ships to satisfy this capacity (Flower 1998 [23], p. 96), this abstraction only disregards the stochasticity of output, due to downtime for example. This tends to be slight (Flower 1998 [23], p. 96). We illustrate the case of identical ships here, but, as discussed below, we can also account for heterogeneity in ship types. The fleet is operated at an average speed of s knots, taken as given in this paper, expressed as nautical miles (NM) per hour. Ships are dedicated and loop between the liquefaction and regasification facilities, which are located at a distance of D NM apart. The average transit time of a ship to sail one-way between these facilities at speed s is D/s, so that the average total *transit* time is $\tau^T(D,s) := 2D/s$. We denote $\tau^L(C)$ and $\tau^U(C)$ the average loading and unloading times (expressed in hours) of a ship with cargo size C at the liquefaction and regasification facilities, respectively. If there were neither uncertainty nor queueing, letting H be the number of operating hours in a year, the capacity of the shipping system would be $\overline{\Lambda} := H \cdot N \cdot C/[\tau^L(C) + \tau^T(D,s) + \tau^U(C)]$. We take $\overline{\Lambda}$ to be the capacity of the liquefaction facility, which is typically measured in MTPA. (We present the conversion factors between the relevant pairs of units of measurement in §3.4.)

To account for uncertainty in shipping operations, we model the loading and unloading blocks as first-come-first-served (FCFS) exponential queues, and the transit blocks as ample-server (AS) exponential or multistage Coxian queues, which allows more flexibility in modeling variability than the exponential distribution (Osogami and Harchol-Balter 2006

[46]). The loading and unloading blocks concisely model the following activities: entering the port (traversing the entry channel), loading/unloading the ship, and leaving the port (traversing the exit channel or again the entry channel if there is only one channel at the given port). Here, randomness in service times does not typically arise when physically loading/unloading the ships, but rather when entering/leaving the port. Uncertainty in transit time is typically due to changing weather conditions.

With this representation, our CQN is a so called multi-class (if there are heterogeneous ship types) "BCMP-network" having closed-form, product-form stationary distributions, as defined and proved by Baskett et al. (1975 [3]). Let I be the total number of blocks (four in Figure 3.2) and J_i the number of Coxian service time stages in block $i = 1, \ldots, I$. The total number of stages is $J = \sum_{i=1}^{I} J_i$. There are K classes, i.e., each type corresponding to a different ship type. In each block i, denote n_{ijk} the number of ships of class $k = 1, \ldots, K$ in stage $j = 1, \ldots, J_i$. The state of the shipping system is array $n = (n_{ijk},\ i = 1, \ldots, I,\ j = 1, \ldots, J_i,\ k = 1, \ldots, K)$. We denote $\mathcal{N}$ the set of all possible states of the system, i.e., the set of states such that $\sum_{i=1}^{I} \sum_{j=1}^{J_i} \sum_{k=1}^{K} n_{ijk} = N$. For each block i, we let λ_{ijk} and μ_{ijk}, respectively, be the total arrival and service rates of stage $j = 1, \ldots, J_i$ for each class k. Baskett et al. (1975 [3]) show that the steady-state probability that the random variable state of the system, $\tilde{n}$, is equal to $n \in \mathcal{N}$, i.e., $\pi(n) := \Pr\{\tilde{n} = n\}$, satisfies the product form solution $\pi(\tilde{n} = n) = \Gamma \prod_{i=1}^{I} \prod_{j=1}^{J_i} \prod_{k=1}^{K} \gamma_{ijk}(\lambda_{ijk}, \mu_{ijk}, n_{ijk})$; here Γ is a normalizing constant chosen to make these probabilities sum to 1, and each term $\gamma_{ijk}(\cdot)$ is computed as follows:

$$\begin{aligned}
\gamma_{ijk}(\lambda_{ijk}, \mu_{ijk}, n_{ijk}) &= \begin{cases} B_{ijk}(\lambda_{ijk}, \mu_{ijk}, n_{ijk}) \left(\sum_{k=1}^{K} n_{ijk}\right)! & \text{if block } i \text{ is FCFS} \\ B_{ijk}(\lambda_{ijk}, \mu_{ijk}, n_{ijk}) & \text{if block } i \text{ is AS} \end{cases} \\
B_{ijk}(\lambda_{ijk}, \mu_{ijk}, n_{ijk}) &:= \prod_{k=1}^{K} \frac{1}{n_{ijk}!} \left(\frac{\lambda_{ijk}}{\mu_{ijk}}\right)^{n_{ijk}}.
\end{aligned}$$

For each class k, define $\mathcal{N}'_k$ the subset of states in which at least one ship is unloading, i.e., $\mathcal{N}'_k := \{n \in \mathcal{N} :\ n_{31k} > 0\}$, where n_{31k} is the number of ships of class k being unloaded at the downstream facility (block 3 in Figure 3.2). Once the probabilities $\pi(\cdot)$'s are known, the throughput of the LNG chain is $\Lambda = H \sum_{k=1}^{K} \mu_{31k} \sum_{n \in \mathcal{N}'_k} \pi(n)$. Since uncertainty in shipping operations can create congestion, it holds that $\Lambda \leqslant \overline{\Lambda}$.

Rolling-forward method. What we have described so far is well-known in the existing literature. Here is how we build on these results to more realistically capture the variability in the unloading distribution. Note that $\pi(n)$ is the stationary distribution for the ship location within the network at a given instant in time. The quantity of interest to us is the random variable $\tilde{u}(\tau)$, the amount of LNG unloaded at the regasification facility by the ships in the CQN during a period of time of length τ, e.g., one month. This is a fundamentally different

quantity from $\pi(n)$; whereas $\pi(n)$ is a distribution at a single instant in time, $\tilde{u}(\tau)$ tracks the distribution of the evolution of the system over an interval of length τ. The distribution of $\tilde{u}(\tau)$ is required in the inventory-release model: for some choice of parameter τ, we need to know the set of its (nonzero-probability) realizations $\mathcal{U}(\tau)$ and the probability $\Pr\{\tilde{u}(\tau) = u\}$ for each $u \in \mathcal{U}(\tau)$. Our rolling-forward method uses the stationary distribution $\pi(\cdot)$ as a *starting point* for this purpose. We allow the CQN to transition forward through time from its stationary distribution, tracking the unloaded amounts and their distribution over a time increment of length τ. To calculate this distribution we uniformize (see, e.g., Asmussen 2003 [1]), condition on the number of "events" that can happen over the time interval of length τ, and then calculate the distribution of the number of ships unloaded (or the unloaded LNG by multiplying this number by the cargo size C) given the number of total events. Given the dynamics of the system, if τ is finite, it is easy to see that set $\mathcal{U}(\tau)$ is finite (provided that the other parameters describing the system take on "sensible" values). Hence, for a given τ, the distribution of random variable $\tilde{u}(\tau)$ is a one-dimensional table. We compute this distribution *analytically*, i.e., we do not have to use Monte Carlo simulation. Denote $\tilde{a}(\tau)$ the random number of events that can occur during a period of time of length τ, a Poisson random variable with opportunely computed mean, and $\tilde{\chi}(\tau)$ the random number of unloaded ships during this time period. Let $\overline{a}$ be such that $\Pr\{\tilde{a}(\tau) > \overline{a}\} = \epsilon$ for arbitrarily small $\epsilon \in \Re_+$. Denote $P(\chi|a, n)$ the probability that χ ships are unloaded during length of time τ given that the system is initially in state n and a total of a events occur. We have developed analytical expressions for $P(\chi|a, n)$ through a forward recursion in χ, but, since they are somewhat lengthy, in the interest of space we do not present them here. We compute the distribution of $\tilde{\chi}(\tau)$, and hence that of $\tilde{u}(\tau) \equiv \tilde{\chi}(\tau)C$, as

$$\Pr\{\tilde{\chi}(\tau) = \chi\} = \frac{\sum_{n\in\mathcal{N}} \pi(n) \sum_{a=0}^{\overline{a}} \Pr\{\tilde{a}(\tau) = a\} P(\chi|a, n)}{1 - \epsilon}.$$

Note that the only inaccuracy in the rolling-forward method is its failure to capture the dependence between unloaded amounts in successive periods of length τ (because we assume the system returns to its stationary distribution $\pi(n)$). As we consider systems with relatively low variability over long horizons (20-30 years), this dependence is not critical.

3.2.2 Inventory-Release Model

Our focus now turns to the optimal management of the downstream buffer and the LNG re-gasification and natural-gas sale activity. We employ a finite-horizon periodic-review model. We define set $\mathcal{T} := \{1, \ldots, T\}$ as the set of time periods, with T a finite integer, and use set $\mathcal{T} \cup \{T + 1\}$ to index the stages of our stochastic dynamic-program. We define time period

$t \in \mathcal{T}$ as the time interval between successive stages t and $t+1$ of length $\overline{\tau}$. We define time period $T+1$ as the last period of interest and denote its length τ_{T+1}, which may be different from $\overline{\tau}$. Typical values for the length of the finite time horizon are 20-30 years, so if we take $\overline{\tau}$ to be one month, then T is typically between 240 and 360.

Unloading distributions. We employ random variable $\tilde{u}_t$ to represent the amount of LNG unloaded at the regasification facility during each time period $t \in \mathcal{T}$. We use the rolling-forward method described in §3.2.1 to construct random variable $\tilde{u}(\overline{\tau})$ and its probability distribution on set $\mathcal{U}(\overline{\tau})$. We define $\tilde{u}_t := \tilde{u}(\overline{\tau})$ and $\mathcal{U}_t := \mathcal{U}(\overline{\tau})$, $\forall t \in \mathcal{T}$. As mentioned previously, this entails the assumption that the random variables in set $\{\tilde{u}_t,\ t \in \mathcal{T}\}$ are independent and identically distributed. This assumption is required in order to make our dynamic-program tractable. Note that we perform this step *before* solving the inventory-release model, which we formulate below.

Denote x_t the inventory available at the downstream facility at the beginning of stage $t \in \mathcal{T} \cup \{T+1\}$. Let 0 and $\overline{x}$, respectively, denote the minimum and maximum levels of inventory that can be held in storage at this facility. Hence, the quantity x_t is constrained to be in set $\mathcal{X} := [0, \overline{x}]$ for all $t \in \mathcal{T} \cup \{T+1\}$. This set poses *space* constraints on the inventory level at the beginning of each stage. The capacity of the downstream facility during a given period of time is the maximum amount of LNG that can be regasified and released during this time period. We denote $\overline{q}_{T+1}$ this capacity in the last time period, and assume that any inventory level $x_{T+1} \in \mathcal{X}$ can be regasified during this period, i.e., $\overline{q}_{T+1} \geqslant \overline{x}$. Since all the other time periods $t \in \mathcal{T}$ share the same length, we denote $\overline{q}$ this capacity for each such time period. Denote q_t the amount of LNG actually regasified and released from downstream storage during time period $t \in \mathcal{T}$. Quantity q_t is constrained to be in set $\mathcal{Q}(x_t, u_t)$, $\forall t \in \mathcal{T}$, $x_t \in \mathcal{X}$, $u_t \in \mathcal{U}_t$. This set poses space *and* capacity constraints on the amount of LNG regasified during any time period. Define $\overline{u}$ the maximum quantity that can be unloaded during any time period, i.e., $\overline{u} := \max\{u : u \in \mathcal{U}(\overline{\tau})\}$ (recall that $\mathcal{U}_t \equiv \mathcal{U}(\overline{\tau})$, $\forall t \in \mathcal{T}$). Given $t \in \mathcal{T}$, $x_t \in \mathcal{X}$, $u_t \in \mathcal{U}_t$, and quantity q_t, the inventory level at time $t+1$ is $x_{t+1} = x_t + u_t - q_t$. It must hold that $x_{t+1} \in [0, \overline{x}]$ and $q_t \in [0, \overline{q}]$. We make the assumption that $\overline{u} \leqslant \overline{q}$, which means that in each time period it is possible to unload *and* regasify the maximum quantity that could be unloaded in this time period. With this assumption, it holds that $\mathcal{Q}(x_t, u_t) = [\max(0, x_t + u_t - \overline{x}), \min(\overline{q}, x_t + u_t)]$.

Remark 1 (Regasification capacity). The assumption that the regasification capacity in each period $t \in \mathcal{T}$ is at least equal to $\overline{u}$ effectively constrains the fleet composition and/or speed in every period, and means that the LNG chain throughput is Λ. (In contrast we refer to $\overline{\Lambda}$ as its capacity.)

Price dynamics and costs. We use the Markovian stochastic process $\{\tilde{p}_t,\ t \in \mathcal{T} \cup \{T+1\}\}$ to describe the evolution of a price state-random-variable $\tilde{p}_t$ with realizations $p_t \in \mathcal{P}_t \subset \Re$. The natural-gas spot price at time $t \in \mathcal{T} \cup \{T+1\}$ is known function $g_t(p_t) : \mathcal{P}_t \to \Re_+$. This price is expressed in U.S. dollars per million British thermal units ($/MMBTU). This price plays a critical role in our model because we use it to convert the physical flow of LNG into a monetary value: we account for regasification sales during time period $t \in \mathcal{T} \cup \{T+1\}$ by multiplying the released quantity, net of regasification fuel losses (explained below), by the natural-gas price prevailing at the *beginning* of this time period. (Using the natural-gas price at the beginning of a time period is a modeling simplification, which we believe is justified given that the typical length of an LNG project, and hence of the time horizon in our model, is of the order of 20-30 years.) Here, we assume that the quantity sold does not affect the market price of natural gas. We also assume that random variable $\tilde{p}_{t+1}|p_t$ is independent of random variable $\tilde{u}_t$, $\forall t \in \mathcal{T}$, which means that we assume that the amount of LNG unloaded in a period does not influence the wholesale-price of natural gas at the downstream location. While we realize that these two assumptions may not always hold in practice, we point out that they should be fairly realistic for modeling an LNG system whose regasification terminal is located in the southearn part of the U.S., e.g., Louisiana, where the natural-gas spot market is fairly liquid.

Remark 2 (Units of measurement). LNG capacity is typically measured in MTPA, cargos in CM, and LNG downstream storage space and regasification capacity in BCF and BCF/day. However, for ease of exposition and consistency with the units of measurement of the natural-gas price, in the inventory-release model we assume that all the physical quantities are expressed in MMBTU, i.e., unloaded LNG cargos and LNG inventory are measured in MMBTU, and regasified LNG and regasification capacity are measured in MMBTU per (the appropriate) unit of time.

There are operating costs and fuel requirements associated with the physical flows along the chain. Denote $h_t(x_t) : \mathcal{X} \to \Re_+$ the physical inventory-holding cost charged against the inventory x_t available at the beginning of period t at the downstream terminal ($h_t(0) = 0$). We assume this function to be (weakly) convex. Let φ^R the fuel (natural gas) needed to regasify one unit of LNG, i.e., the LNG to natural-gas yield is $1 - \varphi^R$. Finally, denote c_t^U the cost of unloading (handling) one unit of LNG in period t at the downstream terminal, i.e., the per unit cost of receiving a cargo. The operating cost of downstream LNG storage and regasification in time period t is $g_t(p_t)\varphi^R q_t + h_t(x_t) + c_t^U u_t$. We introduce the operating costs of natural-gas liquefaction and LNG shipping in period t after having discussed our stochastic dynamic-program.

Stochastic dynamic-programming formulation. The following material uses financial-economics concepts discussed in Luenberger (1998 [39], Chapters 12 and 16) at an introductory level, and in Duffie (2001 [15], Chapters 2 and 6) at an advanced level. We are interested in computing the value of an LNG chain at the beginning of the evaluation period, i.e., $t = 1$. We assume existence of an arbitrage-free natural-gas futures market at the location of the regasification facility, e.g., the NYMEX natural-gas futures and basis-swap market in North America. Since our time horizon is typically longer than the maturity of traded natural-gas futures contracts, we face an incomplete market setting. Stated differently, the futures market would only be complete up to the time of the expiration of the last traded futures contract, e.g., seventy-two months for the NYMEX natural-gas futures contract. Hence, there exist multiple risk-neutral measures that describe the evolution of the stochastic variable $\tilde{p}_t$, *for all* times $t \in \mathcal{T} \cup \{T+1\}$. We define $E_t^*[\cdot] := E^*[\cdot|\tilde{p}_t = p_t]$ as conditional expectation given the current realization p_t of random variable $\tilde{p}_t$ with respect to a measure in the set of all the possible risk-neutral measures. Since there is no market for the unloading risk represented by random variable $\tilde{u}_t$, $\forall t \in \mathcal{T}$, we also face an incomplete market with respect to this risk. By well-known results (see, e.g., Smith and Nau 1995 [59]), we resolve both these incompletenesses for valuation purposes by assuming that the LNG chain, or a portion of it whenever relevant, is operated by a risk-neutral firm.

We now formulate our stochastic-dynamic-programming valuation model. Denote δ_t the risk-free discount factor from time $t+1$ back to time $t \in \mathcal{T}$. In each stage $t \in \mathcal{T} \cup \{T+1\}$, the state of the system is the pair (x_t, p_t), $\forall x_t \in \mathcal{X}$, $p_t \in \mathcal{P}_t$. The optimal value-function in this state and stage is $V_t(x_t, p_t)$. We intepret this quantity as the optimal downstream operating margin-to-go, i.e., this value is the expected optimal revenue minus the expected storage and regasification operating costs during a period plus the discounted optimal expected-margin from the next period onward (however this quantity does *not* account for the liquefaction and shipping operating costs, which will be included as needed below). This function is defined as

$$\begin{aligned}
V_{T+1}(x_{T+1}, p_{T+1}) &:= g_{T+1}(p_{T+1})(1-\varphi^R)x_{T+1} - h_{T+1}(x_{T+1}), \forall x_{T+1} \in \mathcal{X},\ p_{T+1} \in \mathcal{P}_{T+1} \\
V_t(x_t, p_t) &= E\left[v_t\left(x_t, p_t, \tilde{u}_t\right)\right],\ \forall t \in \mathcal{T},\ x_t \in \mathcal{X},\ p_t \in \mathcal{P}_t \\
v_t(x_t, p_t, u_t) &:= \max_{q_t \in \mathcal{Q}(x_t, u_t)} \nu_t(x_t, q_t, p_t, u_t) \\
\nu_t(x_t, q_t, p_t, u_t) &:= g_t(p_t)(1-\varphi^R)q_t - h_t(x_t) - c_t^U u_t + \delta_t E_t^*[V_{t+1}(x_t + u_t - q_t, \tilde{p}_{t+1})].
\end{aligned}$$

This formulation is interpreted as follows. In the final time period $T+1$, any remaining inventory is released and sold (recall our assumption that $\overline{q}_{T+1} \geqslant \overline{x}$). During each time period $t \in \mathcal{T}$, an amount $u_t \in \mathcal{U}_t$ of LNG is unloaded at the downstream terminal at cost $c_t^U u_t$, and becomes available for regasification and sale during time period t; holding cost

$h_t(x_t)$ is charged against the initial inventory x_t; an optimal amount $q_t^*(x_t, p_t, u_t)$ of LNG from the total available inventory $x_t + u_t$ is regasified, and a fraction $1 - \varphi^R$ of this amount is sold into the wholesale natural-gas spot market at the prevailing price $g_t(p_t)$. A stochastic transition to the next period is made to account for the uncertainty in the natural-gas spot price, taking into account the inventory dynamics.

Producer and merchant optimal value-functions. We can further specify value functions for the integrated producer and the LNG merchants who purchase LNG from a nonintegrated producer (see the discussion in §3.1). The integrated producer incurs liquefaction and shipping operating costs at the beginning of each stage t. Since we do not model the operational details of natural-gas liquefaction and LNG shipping, we simply denote these costs at the beginning of stage t as OC_t^L and OC_t^S (we illustrate these costs in §3.4). As stated in §3.1, we consider two types of merchants, type 1 is only active in the regasification of LNG, and type 2 is also active in the shipping of LNG. Let $\tilde{\ell}_t$ the random variable amount of LNG loaded during period t. We make the same assumptions on $\tilde{\ell}_t$ that we make on $\tilde{u}_t$, and compute its distribution and support in a fashion analogous to the computation of those for $\tilde{u}_t$ (we have also developed an extended version of the rolling-forward method to compute their joint distribution). Denote $\zeta_t^{M1}(p_t, \ell_t)$ and $\zeta_t^{M2}(p_t, u_t)$ generic time period t LNG price functions for merchants types 1 and 2, respectively, i.e., the price paid by a merchant of either type to purchase LNG from a nonintegrated producer. According to industry practices, these functions depend on the natural gas (or oil) price and the amount purchased (see, e.g., Miller 1998 [44], pp. 172-176). We will impose more structure on these functions in §3.4. Also, notice that we introduce merchant-type specific functions because the type-1 merchant also incurs the shipping operating cost, whereas the type-2 merchant does not. Thus, *ceteris paribus*, the price paid by the type-2 merchant should be higher than that paid by the type-1 merchant because the shipping cost in the former case is incurred by the nonintegrated producer. The (net) optimal value functions for the integrated producer and the two merchant types are

$$
\begin{aligned}
V_{T+1}^P(x_{T+1}, p_{T+1}) &\equiv V_{T+1}^{M1}(x_{T+1}, p_{T+1}) \\
&\equiv V_{T+1}^{M2}(x_{T+1}, p_{T+1}) := V_{T+1}(x_{T+1}, p_{T+1}), \forall x_{T+1} \in \mathcal{X},\ p_{T+1} \in \mathcal{P}_{T+1} \\
V_t^P(x_t, p_t) &:= V_t(x_t, p_t) - \left(\text{OC}_t^L + \text{OC}_t^S\right),\ \forall t \in \mathcal{T},\ x_t \in \mathcal{X},\ p_t \in \mathcal{P}_t \\
V_t^{M1}(x_t, p_t) &:= V_t(x_t, p_t) - \left\{E\left[\zeta_t^{M1}(p_t, \tilde{\ell}_t)\right] + \text{OC}_t^S\right\},\ \forall t \in \mathcal{T},\ x_t \in \mathcal{X},\ p_t \in \mathcal{P}_t \\
V_t^{M2}(x_t, p_t) &:= V_t(x_t, p_t) - E\left[\zeta_t^{M2}(p_t, \tilde{u}_t)\right],\ \forall t \in \mathcal{T},\ x_t \in \mathcal{X},\ p_t \in \mathcal{P}_t.
\end{aligned}
$$

These definitions simply reflect the different roles played by these players.

3.2.3 The Value and Benefit of Downstream Storage

We define the value of the option to store as the additional value obtained by optimally managing the downstream inventory relative to the value of regasifying and selling each incoming cargo upon receipt. It is clear than if we assume that the storage facility is empty at the beginning of the first stage, i.e., $x_1 = 0$, then the inventory at the beginning of any other stage under the latter policy is always zero. We make this assumption in this paper. Thus, we can suppress the inventory level from the state definition of this policy and denote $\underline{V}_t(p_t)$ its value in stage t with realized price state-variable p_t. Hence, our stochastic dynamic-program simplifies to the following Markov reward process:

$$\begin{aligned} \underline{V}_{T+1}(p_{T+1}) &:= V_{T+1}(0, p_{T+1}) = 0, \forall p_{T+1} \in \mathcal{P}_{T+1} \\ \underline{V}_t(p_t) &= [g_t(p_t)(1-\varphi^R) - c_t^U]E[\tilde{u}_t] + \delta_t E_t^*[V_{t+1}(\tilde{p}_{t+1})],\ \forall t \in \mathcal{T},\ p_t \in \mathcal{P}_t. \end{aligned}$$

We introduce functions $\underline{V}_t^P(\cdot)$, $\underline{V}_t^{M1}(\cdot)$, and $\underline{V}_t^{M2}(\cdot)$ by replacing $V_t(0,\cdot)$ with $\underline{V}_t(\cdot)$ in the definitions of functions $V_t^P(\cdot,\cdot)$, $V_t^{M1}(\cdot,\cdot)$, and $V_t^{M2}(\cdot,\cdot)$, respectively.

Assuming that $x_1 = 0$, we define the value of storage in stage t and state (x_t, p_t) as

$$S_t(x_t, p_t) := V_t(x_t, p_t) - \underline{V}_t(p_t),$$

which is clearly nonnegative. We do not include the type of LNG player in this definition because

$$V_t(x_t, p_t) - \underline{V}_t(p_t) \equiv V_t^P(x_t, p_t) - \underline{V}_t^P(p_t) \equiv V_t^{M1}(x_t, p_t) - \underline{V}_t^{M1}(p_t) \equiv V_t^{M2}(x_t, p_t) - \underline{V}_t^{M2}(p_t).$$

On the contrary, the benefit of storage, or, equivalently, the relative value of optimizing the inventory-management policy, *does* depend on the type of LNG player because the integrated producers and the two types of merchants have different operating income streams. We define the benefit of storage for an integrated producer and the two merchant types as follows:

$$\begin{aligned} B_t^P(x_t, p_t) &:= \frac{S_t(x_t, p_t)}{\underline{V}_t^P(p_t)} \equiv \frac{V_t^P(x_t, p_t)}{\underline{V}_t^P(p_t)} - 1 \\ B_t^{M1}(x_t, p_t) &:= \frac{S_t(x_t, p_t)}{\underline{V}_t^{M1}(p_t)} \equiv \frac{V_t^{M1}(x_t, p_t)}{\underline{V}_t^{M1}(p_t)} - 1 \\ B_t^{M2}(x_t, p_t) &:= \frac{S_t(x_t, p_t)}{\underline{V}_t^{M2}(p_t)} \equiv \frac{V_t^{M2}(x_t, p_t)}{\underline{V}_t^{M2}(p_t)} - 1. \end{aligned}$$

3.3 Optimal Inventory-Management

We now study the structure of the optimal policy to manage LNG inventory at the regasification terminal, i.e., we address research question (2). This structure turns out to be a price

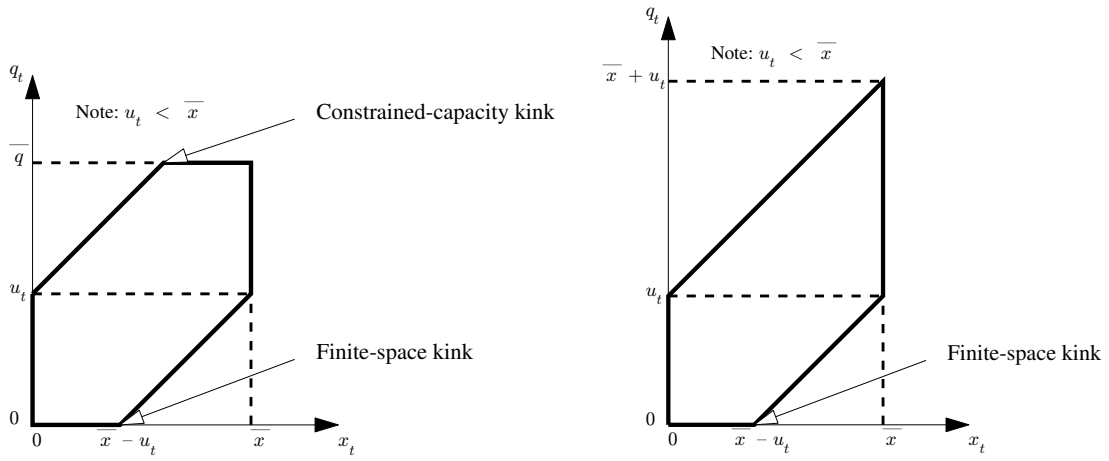

(a) $\mathcal{A}(u_t)$ (b) $\mathcal{A}'(u_t)$

Figure 3.3: The inventory-action spaces $\mathcal{A}(u_t)$ and $\mathcal{A}'(u_t)$ with the roles played by finite capacity and space highlighted.

state-variable dependent basestock target policy. We start by characterizing the optimal value-function in Proposition 5, whose proof is in Appendix 3.6.

Proposition 5 (Optimal value-function). *In any stage $t \in \mathcal{T} \cup \{T+1\}$, it holds that (i) $V_t(x_t, p_t)$ is concave in $x_t \in \mathcal{X}$ for each $p_t \in \mathcal{P}_t$; (ii) if the probability distribution of $\tilde{p}_{t+1}$ conditional on $\tilde{p}_t$ being equal to $p_t \in \mathcal{P}_t$ is stochastically increasing in p_t and $g_t(p_t)$ is increasing in p_t, then $V_t(x_t, p_t)$ is supermodular in (x_t, p_t) on $\mathcal{X} \times \mathcal{P}_t$.*

A few comments on this result are in order. Even though we have assumed a convex physical holding-cost function, the concavity result (i) is fundamentally a consequence of the kinked inventory-action space $\mathcal{A}(u_t) := \{(x_t, q_t) : \ x \in \mathcal{X}, \ q_t \in \mathcal{Q}(x_t, u_t)\}$ illustrated in panel (a) of Figure 3.3. In other words, this property would persist even if one were to set the holding cost function to be identically zero. In addition, this result does not depend on specific distributional assumptions for the stochastic processes $\{\tilde{u}_t, \ t \in \mathcal{T}\}$ and $\{\tilde{p}_t, \ t \in \mathcal{T} \cup \{T+1\}\}$. In point (ii), the assumptions that the probability distribution of $\tilde{p}_{t+1}$ is stochastically increasing in p_t and that function $g_t(p_t)$ is increasing in p_t are not restrictive. For example, they are satisfied when the spot price dynamics follow the mean-reverting model reviewed by Schwartz (1997 [50]) and extended by Jaillet et al. (2004 [31]) to account for deterministic seasonality, or the short-term/long-term model of Schwartz and Smith (2000 [51]). Under these assumptions, the supermodularity of the optimal value-function in the price state-variable and inventory in point (ii) indicates an appealing complementarity relationship between price and inventory: the higher the price state-variable, the more inventory one

would like to have available to sell. Concavity plays a fundamental role in establishing the price state-variable dependent basestock target nature of the optimal inventory-management policy in Proposition 6, and supermodularity allows us to further characterize the behavior of this structure in the price state-variable in Proposition 7.

Basestock target and forced/optional sales. Define $\mathcal{Q}'(x_t, u_t) := [\max(0, x_t + u_t - \overline{x}), x_t + u_t]$, i.e., remove the effect of the capacity limit $\overline{q}$ on set $\mathcal{Q}(x_t, u_t)$, and define the corresponding inventory-action space $\mathcal{A}'(u_t) := \{(x_t, q_t) : \ x \in \mathcal{X}, \ q_t \in \mathcal{Q}'(x_t, u_t)\}$ (see panel (b) of Figure 3.3). We define $q_t^*(x_t, p_t, u_t)$ and $q_t^\diamond(x_t, p_t, u_t)$ as the largest quantities in sets $\arg\max_{q_t \in \mathcal{Q}(x_t, u_t)} \nu(x_t, q_t, p_t, u_t)$ and $\arg\max_{q_t \in \mathcal{Q}'(x_t, u_t)} \nu(x_t, q_t, p_t, u_t)$, respectively.

Proposition 6 (Basestock target). *In any stage $t \in \mathcal{T}$, given a price state-variable realization $p_t \in \mathcal{P}_t$, there exists a basestock level $b_t(p_t) \in \mathcal{X}$ such that in extended state $(x_t, u_t, p_t) \in \mathcal{X} \times \mathcal{U}_t \times \mathcal{P}_t$ (i) ignoring the capacity restriction $q_t \leqslant \overline{q}$, it is optimal to sell down to $b_t(p_t)$ from inventory level $x_t + u_t$, i.e., $q_t^\diamond(x_t, p_t, u_t) = \max(0, x_t + u_t - b_t(p_t))$; (ii) taking the capacity constraint into account, it is optimal to try to sell down to $b_t(p_t)$, i.e., $q_t^*(x_t, p_t, u_t) = \min(\overline{q}_t, q_t^\diamond(x_t, p_t, u_t))$.*

Before discussing how one could prove this result, we note that the quantity $b_t(p_t)$ is a target because the limited regasification capacity $\overline{q}$ may make it unreachable from some initial inventory levels. Also, notice that both $q_t^\diamond(\cdot, \cdot, \cdot)$ and $q_t^*(\cdot, \cdot, \cdot)$ cannot be smaller than $\max(0, x_t + u_t - \overline{x})$, which highlights an important connection between shipping logistics and the management of the regasification terminal in terms of available space: in time period t one *must* release an amount of inventory equal to $\max(0, x_t + u_t - \overline{x})$ to avoid a "tank overflow" due to an incoming shipment. Hence, irrespective of the prevailing natural gas price $g_t(p_t)$, in extended state (x_t, p_t, u_t) one *must* execute *forced* sale $q_t^F(x_t, u_t) := \max(0, x_t + u_t - \overline{x})$. Since the inventory level after performing a *positive* forced sale is $\overline{x}$, this can be intepreted as the forced-sale basestock level (it is not a target because $\overline{q} \geqslant \overline{u}$). We then call *optional* sale the difference between feasible sale q_t and forced sale $q_t^F(x_t, u_t)$ and denote it $\hat{q}_t(x_t, p_t, u_t) := q_t(x_t, p_t, u_t) - q_t^F(x_t, u_t)$. Hence, an optimal sale $q_t^*(x_t, p_t, u_t)$ can be decomposed as $q_t^*(x_t, p_t, u_t) \equiv q_t^F(x_t, u_t) + \hat{q}_t^*(x_t, p_t, u_t)$, and $b_t(p_t)$ assumes the interpretation of basestock target for optimal optional sales. This distinction elucidates the different roles played by forced and optional sales: the former avoids a tank overflow and is fundamentally driven by the incoming shipment u_t, while the latter prevents costly future forced sales or takes advantage of a "high" natural gas price $g_t(p_t)$, and is driven by the residual inventory and capacity after having performed the forced sale.

This analysis also suggests a simplification of the inventory-management problem that facilitates establishing the validity of Proposition 6. Since finding an optimal action is

equivalent to finding an optimal optional sale *after* having performed the forced sale, we denote y_t the post forced-sale inventory level and, given x_t and u_t, define it as

$$y_t := \begin{cases} \overline{x} & x_t + u_t \in (\overline{x}, \overline{x} + \overline{u}] \\ x_t + u_t & x_t + u_t \in [0, \overline{x}] \end{cases}.$$

This inventory level can only take values in set $\mathcal{X}$. By the definition of the function $\nu_t(x_t, q_t, p_t, u_t)$, the costs $h_t(x_t)$ and $c_t^U u_t$ do not affect the choice of an optimal optional sale at time t, and we can restrict our attention for this purpose to state $(y_t, p_t) \in \mathcal{X} \times \mathcal{P}_t$. To find an optimal optional sale in this state, we first consider the relaxed problem of finding an optimal optional sale by ignoring the capacity restriction. Thus, the feasibility set is simply equal to $[0, y_t] =: \check{\mathcal{Q}}_t(y_t)$ for each $y_t \in \mathcal{X}$, and the problem to be solved is $\max_{q_t \in \check{\mathcal{Q}}_t(y_t)} \check{\nu}_t(y_t, q_t, p_t)$ where

$$\check{\nu}_t(y_t, q_t, p_t) := g_t(p_t)(1 - \varphi^R)q_t + \delta_t E_t^*[V_{t+1}(y_t - q_t, \tilde{p}_{t+1})].$$

Consider the capacity constraint on an optional sale $\check{q}_t$ in state (y_t, p_t). If $y_t = \overline{x}$, which occurs if and only if $x_t \in (\overline{x}, \overline{x} + \overline{u}]$, then $\check{q}_t$ must satisfy $\check{q}_t \in [0, \min(\overline{x}, \overline{q} - (x_t + u_t - \overline{x}))]$. This simply says that in this case $\check{q}_t$ must be nonnegative and cannot exceed the minimum of the residual inventory, $\overline{x}$, and the residual capacity after performing forced sale $x_t + u_t - \overline{x}$, $\overline{q} - (x_t + u_t - \overline{x})$. If $y_t \in [0, \overline{x})$, which occurs if an only if $x_t + u_t \in [0, \overline{x}]$, then $\check{q}_t$ must satisfy $\check{q}_t \in [0, \min(y_t, \overline{q})]$. This condition states that $\check{q}_t$ must be nonnegative and cannot be greater than the minimum of the available inventory, y_t, and the regasification capacity, $\overline{q}$ (in this case the forced sale is zero). Thus, an optimal optional sale in state (y_t, p_t) is

$$\check{q}^*(y_t, p_t) := \begin{cases} \min(\check{q}^\diamond(y_t, p_t), \overline{q} - (x_t + u_t - \overline{x})) & y_t = \overline{x} \\ \min(\check{q}^\diamond(y_t, p_t), \overline{q}) & y_t \in [0, \overline{x}) \end{cases}.$$

Since it is also clear that an optimal sale in state (x_t, p_t, u_t) satisfies $q_t^*(x_t, u_t, p_t) \equiv q_t^F(x_t, u_t) + \check{q}_t^*(x_t + u_t - q_t^F(x_t, u_t), p_t)$, an optimal optional sale is $\hat{q}_t^*(x_t, p_t, u_t) \equiv \check{q}_t^*(x_t + u_t - q_t^F(x_t, u_t), p_t)$. The basestock target $b_t(p_t)$ stated in Proposition 6 then satisfies $b_t(p_t) \equiv \overline{x} - \check{q}_t^\diamond(\overline{x}, p_t)$. Given this problem simplification, Proposition 6 can be proved in different ways. One is by adaptation of the method of Secomandi (2007 [53]). Another is by combining results available in Topkis (1998 [61]) with the supermodularity of function $\check{\nu}_t(y_t, q_t, p_t)$ in $(y_t, q_t, p_t) \in \mathcal{X} \times \mathcal{P}_t \times \check{\mathcal{Q}}_t(y_t)$, which, under the stated assumptions, can be established in a manner similar to the proof of part (ii) of Proposition 5 in Appendix 3.6. For completeness, we provide a self-contained proof in Appendix 3.6.

Nontriviality of the basestock target. We now bring to light the nontrivial nature of basestock target $b_t(p_t)$ by addressing the following question: If it is optimal to sell, can it

be optimal to stop selling rather than draining the terminal as much as possible? Formally, we wish to understand if under an optimal policy it can happen that $b_t(p_t) > 0$ in some stage t for some price realization p_t. We show that this can indeed be optimal and that the reasons that cause this phenomenon reside in the interplay among shipping logistics and the operational features of the terminal, i.e., the kinks displayed in Figure 3.3. Example 2 focuses on the role played by the finite-space kink, but it is possible to construct a similar example to show that the constrained-capacity kink can play a similar role.

In Example 2, we take $N = 2$ and, for simplicity, we consider *deterministic* price and shipping dynamics. In particular, we assume that $g_t(p_t) = p_t$, $t = 1, 2, 3$. With a slight abuse of notation, consider prices p_L, p_M, p_H, with $0 < p_L < p_M < p_H$, where subscripts L, M, and H stand for "low," "medium," and "high," respectively. We define the price sets at times 1, 2, and 3 as the singletons $\mathcal{P}_1 := \{p_M\}$, $\mathcal{P}_2 := \{p_L\}$, and $\mathcal{P}_3 := \{p_H\}$. This setting allows us to construct a simplified representation of the marked seasonality displayed by the NYMEX natural-gas forward curve, which is discussed in §3.4. In addition, the amount shipped in each period, u_t, $t = 1, 2$, is deterministic and equal to $\overline{u}$ such that $0 < \overline{u} < \overline{x}$ (recall our convention from §3.2.2 that no shipments occur in stage $N + 1 = 3$). We assume that the discount factor is equal to one in all time periods, $\delta_t = 1$, $t = 1, 2, 3$, and, to emphasize the role played by the space kink, we set all relevant costs and the fuel loss equal to zero: $h_t(\cdot) = 0$ for $t = 1, 2, 3$, $\mathrm{OC}_t^L = \mathrm{OC}_t^S = 0$ and $c_t^U = 0$ for $t = 1, 2$, and $\varphi^R = 0$. The intuition in Example 2 is simple, yet revealing. One would like to sell as much inventory as possible in period 3 to fetch the high price in this period, p_H. However, starting with a full tank in period 1, if one does not release enough inventory in this period, one must sell a positive amount in period 2 at the low price, p_L. Consequently it is optimal to selectively release inventory in period 1, avoid a forced sale in period 2, and sell as much as possible in period 3. This argument is now formally established.

Example 2 (Finite space). Consider a terminal with finite buffer size ($0 < \overline{x} < \infty$) and large regasification capacity ($\overline{q} > \overline{x} + \overline{u}$). As stated, $N + 1 = 3$ and the price at time 1 is at medium level p_M, drops to low level p_L at time 2, and raises to high level p_H at time 3, i.e., $p_1 = p_M$, $p_2 = p_L$, and $p_3 = p_H$. In stage 3, it is clear that $V_3(x_3, p_H) = p_H x_3$ for $x_3 \in [0, \overline{x}]$. In stage 2, for $x_2 \in [0, \overline{x}]$ and $q_2 \in [\max(0, x_2 + \overline{u} - \overline{x}), x_2 + \overline{u}]$, it holds that

$$\nu_2(x_2, q_2, p_L, \overline{u}) = p_L q_2 + V_3(x_2 + \overline{u} - q_2, p_H) = (p_L - p_H)q_2 + p_H(x_2 + \overline{u}).$$

Since $\nu_2(\cdot, \cdot, \cdot, q_2)$ decreases in q_2, it follows that $q_2^*(x_2, p_L, \overline{u}) = q_2^F(x_2, \overline{u}) = \max(0, x_2 + \overline{u} - \overline{x})$, and $V_2(x_2, p_L) = (p_L - p_H)\max(0, x_2 + \overline{u} - \overline{x}) + p_H(x_2 + \overline{u})$. In stage 1, consider inventory

level $x_1 = \overline{x}$. For $q_1 \in [\overline{u}, \overline{x} + \overline{u}]$, it holds that

$$\begin{aligned}
\nu_1(\overline{x}, q_1, p_M, \overline{u}) &= p_M q_1 + V_2(\overline{x} + \overline{u} - q_1, p_L) \\
&= p_M q_1 + (p_L - p_H)\max(0, 2\overline{u} - q_1) + p_H(\overline{x} + 2\overline{u} - q_1) \\
&= \begin{cases} (p_M - p_L)q_1 + 2p_L\overline{u} + p_H\overline{x} & q_1 \in [\overline{u}, 2\overline{u}] \\ (p_M - p_H)q_1 + p_H(\overline{x} + 2\overline{u}) & q_1 \in [2\overline{u}, \overline{x} + \overline{u}] \end{cases}.
\end{aligned}$$

Therefore, $\nu_1(\overline{x}, q_1, p_M, \overline{u})$ reaches a peak at $q_1 = 2\overline{u}$, and $q_1^{\diamond}(\overline{x}, p_M, \overline{u}) = 2\overline{u} < \overline{x} + \overline{u} < \bar{q}$, which shows that with a full terminal in period 1 it is unconstrained optimal to sell down to basestock target $b_1(p_M) = \overline{x} - \overline{u} > 0$. Also notice that a forced sale is avoided in period 2: $q_2^F(\overline{x} - \overline{u}, \overline{u}) = 0$.

Basestock target behavior. We now study the behavior of $b_t(p_t)$ in p_t.

Proposition 7 (Basestock target and price state-variable). *In any stage $t \in \mathcal{T}$, if the probability distribution of $\tilde{p}_{t+1}$ conditional on $\tilde{p}_t$ being equal to $p_t \in \mathcal{P}_t$ is stochastically increasing in p_t and $g_t(p_t)$ increases in p_t, then basestock target $b_t(p_t)$ decreases in p_t.*

This result follows from Theorem 2.8.2 in Topkis (1998 [61], p. 77) and the supermodularity of $\check{\nu}_t(y_t, q_t, p_t)$ in $(y_t, q_t, p_t) \in \mathcal{X} \times \mathcal{P}_t \times \check{\mathcal{Q}}_t(y_t)$, which can be proved in the manner previously stated. It says that as the price state-variable increases it is optimal to release and sell more inventory. Thus, the extended inventory, $x_t + u_t$, and price state-variable, p_t, space is partitioned into two regions, one where it is (unconstrained) optimal to *sell* down to the basestock target, the other where it is optimal to *hold* inventory. Example 3 illustrates this result in a deterministic setting.

Example 3 (Basestock target in Example 2). Consider stage 1 in Example 2. Suppose now that price p_1 can take values in set $[\underline{p}, \overline{p}]$, with $0 < \underline{p} < p_L$ and $p_H < \overline{p} < \infty$. It is easy to see that for $p_1 \in [\underline{p}, p_L]$ the basestock target is $\overline{x}$, for $p_1 \in (p_L, p_H)$ the basestock target is $\overline{x} - \overline{u}$, and for $p_1 \in [p_H, \overline{p}]$ it is zero. Thus, $b_1(p_1)$ decreases in p_1 as displayed in Figure 3.4 (note that $\hat{q}_1^*(\cdot,\cdot,\cdot) \equiv \hat{q}_1^{\diamond}(\cdot,\cdot,\cdot)$). This figure also shows the "hold" and "sell" regions.

3.4 Quantification of the Value and Benefit of Storage

We now report the results of a model-based numerical study of the value and relative benefit of downstream LNG storage conducted to address research questions (3) and (4). After illustrating the study set-up, the natural-gas price model used, and its parameter estimation, we discuss the valuation and benefit assessment results, along with the managerial insights

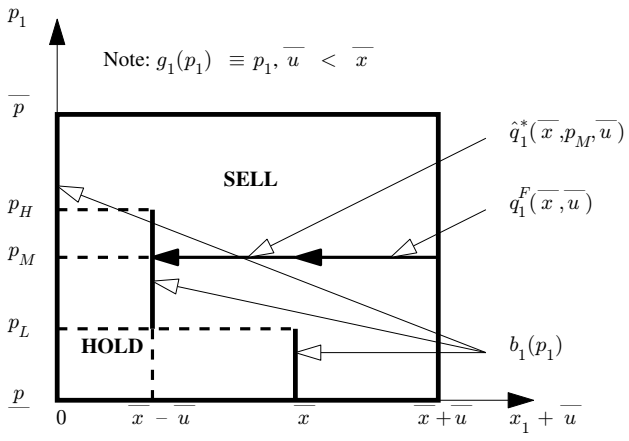

Figure 3.4: The structure of the basestock target in Example 3; note that $\hat{q}_1^\diamond(x_1, p_1, \overline{u}) \equiv \hat{q}_2^*(x_1, p_1, \overline{u})$.

Table 3.1: Units of measurement and conversion factors.

MTPA	Million Tons per Annum
CM	Cubic Meters
NM	Nautical Miles
MMBTU	Million British Thermal Units
BCF	Billion Cubic Feet
1MTPA = 51,982,370MMBTU per annum	
1Knot = 1NM per Hour	
1MMBTU = 23.6863CM	
1BCF = 1,100,000MMBTU	

that follow from them. Table 3.1 reports relevant units of measurements and conversion factors employed in this study.

Operational parameters and operating costs. Table 3.2 summarizes the operational parameters and operating costs of the liquefaction, shipping, and regasification stages used in our numerical experiments. Based on figures reported by Denny (2005 [14]), we take the liquefaction-capacity cost to be \$200M per MTPA, and compute the operating cost per year of the liquefaction plant as 4% of this capacity cost (Flower 1998 [23], p. 95, reports that this percentage is between 3 and 5). We consider a fleet of homogeneous ships, each with a cargo size of 145,000CM, which is a common size in the LNG industry. We compute the ship operating cost to be \$11.2M/year based on the figures reported by Cho et al. (2005 [10]; see Appendix 3.6), which is roughly consistent with the \$12M/year figure reported by Flower (1998 [23], p. 101). We set the distance between the liquefaction and regasification facilities equal to 7,000NM, which is roughly equal to the distance between Egypt, an LNG exporting country, and Lake Charles, Louisiana, which hosts an LNG terminal operated by Trunkline LNG. We assume that the speed of each ship is 19 knots, a realistic value (Flower 1998 [23], p. 100, and Cho et al. 2005 [10]), which makes a one-way trip approximately 15 days long. The mean service times at the liquefaction and regasification facilities are one day each, which are representative of typical operations (EIA 2001 [17]). The unloading (handling) charge and the regasification fuel-loss reflect the "Currently Effective Rates," as of 6/26/2006, for firm terminal service of Trunkline LNG. The downstream storage space is 6BCF, which amounts to two cargos (according to EIA 2001 [17], two cargos is the industry rule of thumb for the size of the receiving tanks). Since the Lake Charles terminal has 9BCF of storage space and some of the new terminals under development in the U.S. have even larger sizes, we also consider larger values for this parameter in our analysis. The send-out capacity is 1BCF/day, which is consistent with the capacity of the Lake Charles and other terminals in the U.S. Apparently, Trunkline LNG, as well as the other companies that manage the active regasification terminals in the U.S., do not charge a holding cost, so we set this to zero in our experiments.

Table 3.3 reports the capacity, throughput, and capacity loss for different fleet sizes when the loading/unloading and transit times are exponentially distributed. (As discussed below, we also experimented with Coxian transit times.) In computing these figures, we assume that ships are operated 365 days per year, 24 hours per day. The capacity losses are minimal, which indicates little congestion in this system, or, equivalently, short ship waiting times at the loading and unloading facilities. However, one should not conclude that the system operations do not exhibit stochastic variability. Table 3.4 displays the probability mass functions of the number of unloaded cargos per month computed by the rolling-forward

Table 3.2: Operational parameters and operating costs.

Liquefaction				
Loading Time				Operating Cost
1 Day				\$8M/MTPA
Shipping				
Transit Time	Distance	Speed	Ship Size	Operating Cost
15 Days	7,000NM	19 Knots	145,000CM	\$11.2M/(Year, Ship)
Regasification				
		Variable Costs		
Unloading Time	Capacity	Fuel Loss	Unloading	Holding
1 Day	1 BCF/Day	1.69%	\$0.0285/MMBTU	-

Note: Processing times exponentially distributed (means displayed); transit time is one way.

Table 3.3: Capacity and throughput for different fleet sizes.

# of Ships	Capacity (MTPA)	Throughput (MTPA)	% Capacity Loss
1	0.74	0.74	0.00
2	1.49	1.48	0.67
3	2.23	2.22	0.45
4	2.97	2.95	0.67
5	3.72	3.68	1.08
6	4.46	4.39	1.57

Table 3.4: Probability mass functions of the number of unloaded cargos in one month.

	# of Ships in the Fleet					
# of Unloaded Cargos	1	2	3	4	5	6
0	0.2811	0.0791	0.0223	0.0063	0.0018	0.0005
1	0.5220	0.2937	0.1242	0.0468	0.0166	0.0057
2	0.1777	0.3728	0.2728	0.1462	0.0672	0.0283
3	0.0191	0.1957	0.3038	0.2514	0.1565	0.0832
4	0.0000	0.0513	0.1878	0.2627	0.2333	0.1615
5	0.0000	0.0074	0.0699	0.1762	0.2346	0.2189
6	0.0000	0.0000	0.0166	0.0792	0.1650	0.2144
7	0.0000	0.0000	0.0026	0.0248	0.0834	0.1554
8	0.0000	0.0000	0.0000	0.0055	0.0310	0.0850
9	0.0000	0.0000	0.0000	0.0009	0.0087	0.0356
10	0.0000	0.0000	0.0000	0.0000	0.0018	0.0116
Mean	0.9348	1.8685	2.7995	3.7259	4.6454	5.5426
Standard Deviation	0.7288	1.0297	1.2605	1.4528	1.6161	1.7376
Coefficient of Variation	0.7796	0.5511	0.4503	0.3899	0.3479	0.3135

method for different fleet sizes. The coefficients of variation displayed at the bottom of this table indicate that there is significant variability in the system.

Natural-gas price model and its estimation. In our study, we assume that regasified LNG is sold into the Louisiana wholesale natural-gas spot market at the Henry-Hub price. We model the evolution of this price as a one-factor mean-reverting process with deterministic seasonality (Jaillet et al. 2004 [31]). Assuming a seasonality cycle of length one, the seasonality factors f_w, $w \in [0,1]$, must satisfy the normalization condition $\int_0^1 \ln f_w dw = 0$. The natural-gas spot price at time t is the product of a deterministic seasonality factor and a stochastic deseasonalized spot price factor: $g_t(\tilde{p}_t) := f_{w(t)}\tilde{p}_t$, with $f_{w(t)}$ the time-t seasonality factor. With a convenient abuse of notation t, the logarithm of the deseasonalized spot price, $X_t := \ln p_t$, evolves as the Ornstein-Uhlenbeck process

$$dX_t = \kappa(\xi - X_t)dt + \sigma dZ_t. \tag{3.1}$$

Hence, starting from any given initial value X_0, X_t reverts over time to long-term level ξ at speed κ, and is subject to random shocks, with volatility σ, driven by standard Brownian motion Z_t. As in Jaillet et al. (2004 [31]), we assume that Z_t is a Brownian motion under a risk-neutral probability measure (as discussed in §3.2.2, uniqueness of this measure is not guaranteed in our setting), which means that the term ξ is already expressed in risk-adjusted

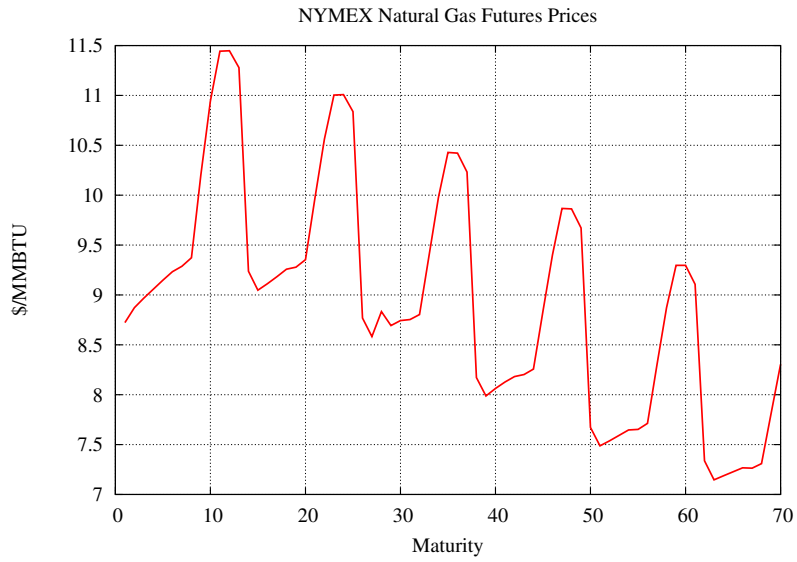

Figure 3.5: NYMEX natural-gas futures prices on 2/1/2006.

form (see Ross 1997 [48], Schwartz 1994 [49], 1997 [50], Schwartz and Smith 2000 [51], and Smith 2005 [56]).

As in Jaillet et al. (2004 [31]) we need to estimate fifteen parameters: κ, ξ, σ, and twelve monthly seasonality factors, $f_1, \ldots, f_{12}$. We use NYMEX data for this purpose and employ a dataset that includes natural-gas futures prices and prices of call and put options on natural-gas futures covering the period from 2/1/2006 to 2/23/2006. Differently from the dataset used by Jaillet et al. (2004 [31]), we have many more option prices at our disposal, so we employ a different estimation approach, which we now describe. Denote $F(t,t')$ the time t price of a futures contract for delivery at time $t' \geqslant t$. It is well known that $F(t,t') = E_t^*[g_{t'}(\tilde{p}_{t'})]$, so that under model (3.1), the following expression for $\ln F(t,t')$ holds (see Jaillet et al. 2004 [31] for details):

$$\ln F(t,t') = \ln f_{w(t')} + e^{-\kappa(t'-t)} X_t + \xi[1 - e^{-\kappa(t'-t)}] + \frac{\sigma^2}{4\kappa}[1 - e^{-\kappa(t'-t)}]. \tag{3.2}$$

By setting $t = 0$, this expression can be used to fit the price model to observed futures prices. However, loosely speaking, futures prices only give an indication of average price behavior and do not allow us to obtain a meaningful estimate of the volatility parameter σ. Option prices can be used to estimate this parameter. Closed-form expressions for the time-0 prices of European call and put options on futures price $F(t_1,t_2)$, $0 < t_1 < t_2$, under model (3.1) depend on the variance term (see Jaillet et al. 2004 [31])

$$\hat{\sigma}^2(t_1,t_2) = e^{-2\kappa(t_2-t_1)}(1 - e^{-2\kappa t_1})\frac{\sigma^2}{2\kappa}. \tag{3.3}$$

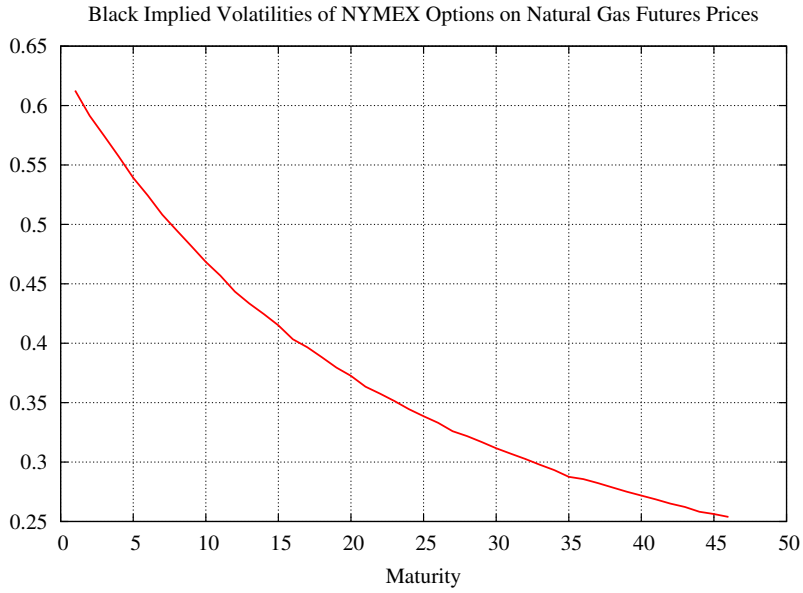

Figure 3.6: Black implied volatilities of NYMEX options on natural-gas futures prices on 2/1/2006.

Given the current (time 0) market price of a traded (call or put) European option on a futures and the futures price, one can use the well-known Black (1976 [4]) formulas for the option price to compute, by means of standard techniques, the so called implied volatility parameter $\hat{\sigma}_B(t_1, t_2)/\sqrt{t_1}$, where subscript B stands for Black. We numerically compute (imply out) a Black's volatility for each option in our dataset. Since under model (3.1) the price of a European call/put option on a futures price would match the market price of the traded option if $\hat{\sigma}_B^2(t_1, t_2)/t_1 = \hat{\sigma}^2(t_1, t_2)$, the implied volatilities and expressions (3.3) provide us with an additional set of conditions to estimate the parameter σ.

Figures 3.5 and 3.6 illustrate the futures prices and Black implied volatilities in our dataset on 2/1/2006 (we have a similar set of data for each trading day of February 2006). The marked seasonality in futures prices is worth noticing. The decline of the volatilities with increasing futures price maturity is typical. For each day in our dataset, similarly to Clewlow and Strickland (2000, p. 160), we compute daily estimates of the price-model parameters by minimizing the sum of the squared *percentage* deviations of observed futures prices and implied variances subject to nonnegativity constraints on κ and σ, and the constraint $\sum_{r=1}^{12} \ln f_r = 0$ (this is a discretized version of the normalization condition $\int_0^1 \ln f_w dw = 0$ that assumes a yearly cycle with monthly seasonality factors that are constant within each month). We then compute estimates of the model parameters by averaging each daily estimate. Table 3.5 displays the parameter estimates and the square root of their mean squared errors. Compared to the estimates of Jaillet et al. (2004), the most significant differences are

Table 3.5: Estimates of the natural-gas price model parameters using NYMEX data from 2/1/2006 to 2/23/2006.

Parameter	Average	Square Root of Mean Squared Error
Long-term log level (ξ)	1.9740	0.0460
Speed of mean reversion (κ)	0.7200	0.1400
Volatility (σ)	0.6610	0.0550
January factor (f_1)	1.1932	0.0157
February factor (f_2)	1.1948	0.0155
March factor (f_3)	1.1412	0.0495
April factor (f_4)	0.9142	0.0114
May factor (f_5)	0.8984	0.0106
June factor (f_6)	0.8998	0.0181
July factor (f_7)	0.9113	0.0159
August factor (f_8)	0.9219	0.0135
September factor (f_9)	0.9287	0.0098
October factor (f_{10})	0.9412	0.0063
November factor (f_{11})	1.0252	0.0034
December factor (f_{12})	1.1049	0.0093

our higher estimate of the long-term level (1.974 vs. 0.802) and lower estimate of the speed of mean reversion (0.72 vs. 3.4), which are due to basic differences in the forward curves between 1997-98 and 2006, the years covered by the data used by Jaillet et al. 2004 and us, respectively. The estimates of the other parameters seem consistent with each others. Our mean squared errors are roughly similar to the standard deviations reported by Jaillet et al. (2004), with the exception of that of the speed of mean reversion, which is much lower in our case.

Storage valuation. We employ a valuation period of twenty years with monthly intervals starting on 3/1/2006, i.e., $t = 1$ corresponds to 3/1/2006. This is consistent with industry practices regarding the valuation of LNG projects (Flower 1998). We discount using a constant annual risk-free rate equal to 5%. We assume that no inventory is available in the tank at $t = 1$, i.e., $x_1 = 0$, and set $g_1(p_1)$ equal to \$8.95/MMBTU, the realized Henry-Hub natural-gas spot price on 3/1/2006 (note that $p_1 = g_1(p_1)/f_3 = \$8.95/1.1412 = \$7.84$, where f_3 is the March seasonality factor). We use the parameter estimates shown in Table 3.5 to build a trinomial lattice that represents the evolution of the deseasonalized natural-gas

Table 3.6: The value of the option to store, $S_1(0, \$7.84)$ ($M; 1 Cargo = 145,000CM = 3BCF).

		Storage Size (# of Cargos)				
# of Ships	Throughput (MTPA)	1	2	3	4	5
1	0.74	*104*	199	284	359	424
2	1.48	107	210	308	400	486
3	2.22	107	213	316	415	509
4	2.95	108	215	320	423	522
5	3.68	108	215	321	425	**526**
6	4.39	107	214	320	422	521

Minimum, **maximum**.

spot-price during the valuation period. Given the length of this period, we use monthly time steps, i.e., $\mathcal{T} := \{1, 2 \ldots, 240\}$. We construct this lattice using standard techniques (see Jaillet et al. 2004 and references therein). However, in the calibration step, we employ the average of the forward-curve values observed in February 2006 for the first 70 months, and we extend this curve to the remaining 170 months, for which prices are not observed, by modifying the latest observed average deseasonalized price according to the seasonality factor corresponding to each maturity. This lattice is the stochastic process $\{p_t,\ t \in \mathcal{T}\}$ used in our inventory-release model. For computational ease, we implemented this model with an inventory state space expressed in terms of number of cargos, which means that in our model implementation the inventory-release actions are multiples of the cargo size. Since this is not necessarily optimal, our estimates of the value of optimizing the inventory-management policy are conservative.

Throughput and space effects. Table 3.6 displays the value of the option to store at time 1, i.e., the quantity $S_1(0, \$7.84)$. This table employs the throughput levels reported in Table 3.3 and different levels of storage space. We only consider system configurations such that the regasification capacity does not impose a constraint on the unloading ships. In other words, we vary the number of ships, while keeping all the other parameters constant, in a manner that ensures that our assumption on the terminal regasification capacity made in §3.2.2 holds, i.e., $\Pr\{\tilde{u}_t > \overline{q}\} = 0$, $\forall t \in \mathcal{T}$. (To be precise, with 4, 5, and 6 ships this statement is true provided that we allow 31 days in each month.) The value of storage varies from $104M (1 ship and 1 cargo of storage space) to $526M (5 ships and 5 cargos of storage space). Table 3.6 illustrates that the value of storage increases in the available space. This is intuitive, as more available space allows more effective exploitation of high natural-

Table 3.7: The extrinsic value of the option to store (% of option value, $S_1(0, \$7.84)$; 1 Cargo = 145,000CM = 3BCF).

		Storage Size (# of Cargos)				
# of Ships	Throughput (MTPA)	1	2	3	4	5
1	0.74	23.9	21.6	19.4	18.1	18.6
2	1.48	24.9	24.3	23.1	21.7	20.5
3	2.22	25.2	24.9	24.4	23.6	22.6
4	2.95	25.2	25.1	24.9	24.5	23.8
5	3.68	25.2	25.2	25.0	24.6	24.1
6	4.39	25.2	25.0	24.7	24.2	23.7

gas prices, i.e., $V_1(0, \$7.84)$ increases while $\underline{V}_1(\$7.84)$ remains constant as the storage space increases. In addition, as one would expect, the marginal benefit of one additional cargo worth of space decreases in storage size as each additional unit of space is used less often. More interestingly, Table 3.6 also shows that the value of storage is nonmonotonic in the number of ships, or, equivalently, throughput. As throughput increases, the total amount of regasified LNG increases and both $V_1(0, \$7.84)$ and $\underline{V}_1(\$7.84)$ increase. But, due to the regasification capacity constraint, these values grow closer to each other after a critical level of throughput, after which the value of storage decreases (see Appendix 3.6 for an intuitive justification of this statement). While the drop in storage value in going from 5 to 6 ships is either none or minimal, the basic insight here is that discretionary regasification capacity must be available for an LNG terminal to have storage value. Otherwise, when throughput approaches the regasification capacity, one must agree with Holcomb (2006) that an LNG terminal is a delivery mechanism with little or no storage value.

Extrinsic value. The extrinsic value of the option to store is the amount of option value that can be attributed to price volatility. We compute this value as follows. We first obtain the intrinsic value of storage by optimizing the inventory-management policy under the assumption that the initial forward curve remains constant over time (the zero volatility case). Then, we compute the extrinsic value by subtracting the intrinsic value from the total option value. Table 3.7 displays the extrinsic value as a percentage of the total option value. Interestingly, this value decreases in storage space for all throughput levels, except when there is only one ship in the fleet and the storage space changes from four to five cargos worth of space. This behavior of the extrinsic value is nonintuitive but can be explained as follows. Increasing the available storage space amounts to increasing the "number" of available (storage) options. As this number grows, the additional space can

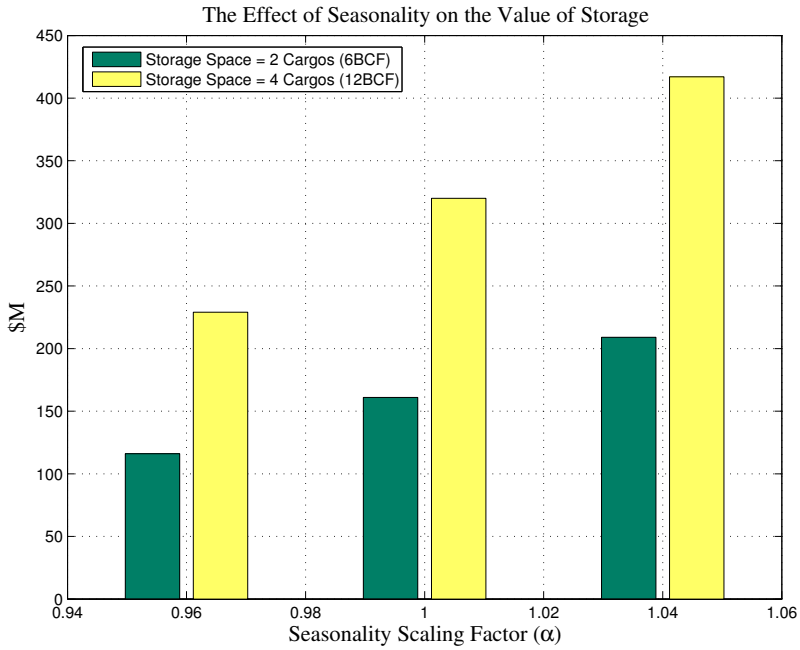

Figure 3.7: Intrinsic value of the option to store at different seasonality levels (6-ship fleet).

be more advantageously used to exploit the seasonality in the natural-gas forward curve. Since this is a deterministic component of the price dynamics (i.e., the seasonality factors are deterministic), the additional value of exploiting it is reflected in the increased intrinsic value of storage.

Seasonality effect. In the natural-gas industry in North America the November-March period is referred to as the heating season, and these months are called the winter months. We assess the effect of seasonality by examining the impact on the value of the option to store of stronger or milder heating seasons, which we achieve by changing the price seasonality factors of these months. In our setting, the winter months have seasonality factors greater than one (factors f_1, f_2, f_3, f_{11}, and f_{12} in Table 3.5). We adjust the magnitudes of these factors by multiplying them by scaling parameter α. Given a value of α, the scaling parameter of the non-winter months, β, is then implied by the normalization condition $\sum_{r=1}^{12} \ln f_r = 0$, i.e., $\beta = \alpha^{-5/(12-5)}$. In addition to the base case of $\alpha = \beta = 1.00$, we consider two possible values for α, 1.04 and 0.96, with corresponding β values equal to 0.9724 and 1.0296.

To emphasize the effect of seasonality (deterministic variability), Figure 3.7 displays the *intrinsic* value of the option to store for two different storage sizes at these three seasonality levels with a 6-ship fleet, i.e., $\Lambda = 4.39$MTPA. With a strong heating season ($\alpha = 1.04$ and $\beta = 0.9724$), the option to store appreciates, intrinsically, by about 30% for each of the two levels of storage space. With a mild heating season ($\alpha = 0.96$ and $\beta = 1.0296$), the option

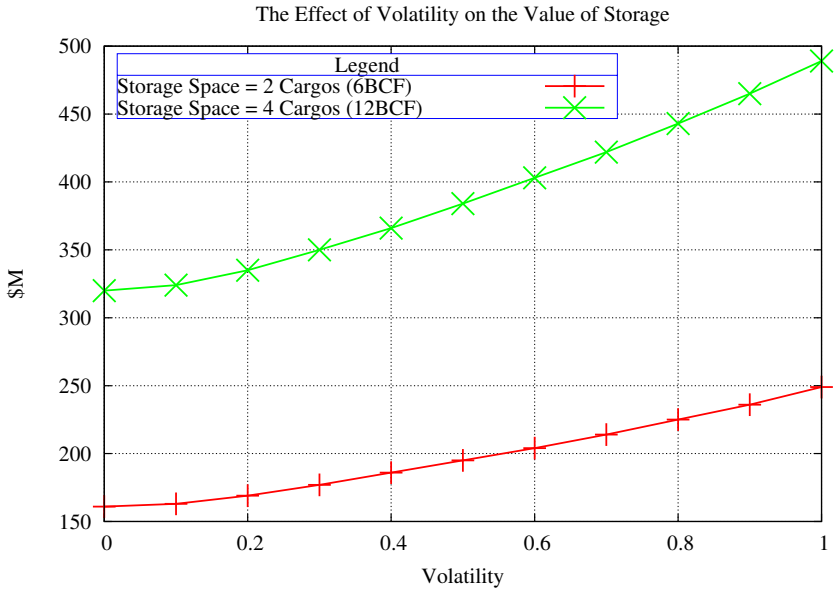

Figure 3.8: The value of the option to store at different volatility levels (6-ship fleet).

depreciates by about 28%. These changes reflect the fact that when the seasonality factors of the winter months increase/decrease, those of the remaining months simultaneously decrease/increase, i.e., the amplitude of the seasonal price oscillations, as displayed for example in Figure 3.5, increases/decreases. We also observe that the seasonality effect is only slightly more pronounced for the larger storage space.

Volatility effect. Figure 3.8 displays the value of the option to store at two storage space levels when changing volatility from 0.00 to 1.00 and keeping seasonality at its base level ($\alpha = \beta = 1.00$; when $\sigma = 0$ we display the intrinsic value of this option) with a 6-ship fleet, i.e., $\Lambda = 4.39$ MTPA. As expected, increasing volatility increases the value of storage. Taking 70% as the base volatility value, changing volatility can add/subtract about 16/25% in value over the range of considered values. Thus, optimal inventory-management becomes more important in a more volatile wholesale natural-gas market, in which case fully capturing the value of the option to store requires making inventory-release decisions based on a stochastic optimization model.

Shipping time variability effect. Exponential transit times simplify the computations of the relevant quantities associated with the shipping model, but a coefficient of variation (CV) equal to 1 is unrealistic for shipping times whose averages are of the order of two or more weeks. Since our CQN model is a BCMP network, we can readily reduce the variability of the transit times by increasing the number of Erlang stages in the two shipping blocks (exponential times correspond to the case of a single Erlang stage). Table 3.8 displays the effect on the value of storage of reducing the variability in shipping times by varying

Table 3.8: The relative effect of reduced shipping time variability on the value of the option to store (displayed values are proportions relative to the case of exponential transit times, i.e., 1 Erlang stage, CV = 1; storage size = 2 cargos.)

	(# of Erlang Stages, CV)			
# of Ships	$(1, 1)$	$(2, 0.7071)$	$(3, 0.5773)$	$(4, 0.5)$
1	100	101.43	102.03	102.39
2	100	100.75	101.03	101.17
3	100	100.38	100.50	100.57
4	100	100.18	100.22	100.25
5	100	100.13	100.15	100.16
6	100	100.29	100.37	100.41

CV: Coefficient of variation of transit times.

Table 3.9: Summary statistics of the unloading distributions at different levels of shipping time variability.

	Mean				CV			
	# of Coxian Stages				# of Coxian Stages			
# of ships	1	2	3	4	1	2	3	4
1	0.9348	0.9374	0.9375	0.9375	0.7796	0.6454	0.5826	0.5434
2	1.8685	1.8712	1.8713	1.8713	0.5511	0.4565	0.4123	0.3848
3	2.7995	2.8011	2.8012	2.8012	0.4503	0.3728	0.3370	0.3146
4	3.7259	3.7267	3.7268	3.7268	0.3899	0.3230	0.2922	0.2730
5	4.6454	4.6476	4.6477	4.6477	0.3479	0.2890	0.2617	0.2447
6	5.5426	5.5613	5.5631	5.5635	0.3135	0.2634	0.2391	0.2238

Table 3.10: The benefit of the option to store for the integrated producer, $B_1^P(0, \$7.84)$; displayed values are percentages.

		Storage Size (# of Cargos)				
# of Ships	Throughput (MTPA)	1	2	3	4	5
1	0.74	2.77	5.31	7.57	9.57	11.28
2	1.48	1.42	2.80	4.10	5.33	6.48
3	2.22	0.95	1.90	2.81	3.69	4.53
4	2.95	0.72	1.43	2.14	2.82	3.49
5	3.68	0.58	1.15	1.72	2.28	2.82
6	4.41	0.48	0.96	1.44	1.90	2.34

Table 3.11: The benefit of the option to store for the type-1 (FOB) merchant, $B_1^{M1}(0, \$7.84)$, with 2 cargos (6BCF) of storage space; displayed values are percentages.

		η^{M1}				
# of Ships	Throughput (MTPA)	0.70	0.75	0.80	0.85	0.90
1	0.74	17.23	20.84	26.36	35.86	56.05
2	1.48	9.89	12.19	15.88	22.77	40.22
3	2.22	6.93	8.60	11.35	16.66	31.34
4	2.95	5.33	6.64	8.82	13.13	25.68
5	3.68	4.32	5.41	7.22	10.84	21.76
6	4.41	3.64	4.56	6.11	9.24	18.94

the number of Erlang stages from one to four. Decreasing this variability increases the value of storage, because, while the mean number of unloaded cargos remains essentially the same, which makes this comparison meaningful, the reduced variability in the number of incoming cargos (Table 3.9) allows for an easier planning of released inventory. However, this increase in value is marginal, or, equivalently, the loss in storage value caused by assuming exponential transit times is very small, especially at higher throughput levels. These findings have practical significance since they suggest that a simple version of the shipping model appears to be adequate for the purposes of storage valuation.

Storage benefit. We now discuss the benefit of inventory management. As in §3.2.3, we consider an integrated producer who manages the entire chain, a type-1 merchant who manages the shipping and regasification parts of the chain, and a type-2 merchant who only manages the regasification stage. Table 3.10 indicates that the benefit of (downstream) storage for the integrated producer are fairly modest, especially at low levels of storage space.

Table 3.12: The benefit of storage for the type-2 (ex-ship) merchant, $B_1^{M2}(0, \$7.84)$, with 2 cargos (6BCF) of storage space; displayed values are percentages.

		η^{M2}				
# of Ships	Throughput (MTPA)	0.75	0.80	0.85	0.90	0.95
1	0.74	18.19	22.26	28.68	40.29	67.69
2	1.48	10.50	13.12	17.50	26.24	52.49
3	2.22	7.36	9.29	12.57	19.43	42.82
4	2.95	5.67	7.18	9.80	15.41	36.14
5	3.68	4.60	5.85	8.02	12.77	31.24
6	4.41	3.88	4.94	6.79	10.90	27.52

This is due to the fact that the value of the flow of LNG is much larger than the value of storage. Nevertheless, these percentage improvements cannot be ignored. For simplicity, in the two merchant cases we specify the following LNG price functions: $\zeta_t^{M1}(p_t, \ell_t) = \eta^{M1} p_t \ell_t$ and $\zeta_t^{M2}(p_t, u_t) = \eta^{M2} p_t u_t$, where η^{M1}, $\eta^{M2} \in (0, 1)$ (these types of functions can be used in practice; see, e.g., Flower 1998, pp. 172-176). This means that the LNG is priced as a fraction of the natural gas price at the downstream location. In addition, it is natural to expect that $\eta^{M1} < \eta^{M2}$, since the type-2 merchant does not incur the shipping cost, and the LNG price paid by this merchant includes the LNG shipping cost incurred by the LNG seller. Hence, we consider the following values for these parameters: $\eta^{M1} \in \{0.70, 0.75, 0.80, 0.85, 0.90\}$ and $\eta^{M2} \in \{0.75, 0.80, 0.85, 0.90, 0.95\}$. Tables 3.11 and 3.12 present the percentage benefit of storage for the two types of merchants with 2 cargo of available space (6BCF). Compared to the integrated producer case, these values are significantly higher, and they are more so for the type-2 merchant. Also, the benefit of storage increases as the price of LNG increases, i.e., when η^{M1} and η^{M2} increase. This benefit quantification shows that optimizing the management of the inventory held at the regasification facility is of paramount importance for merchants.

3.5 Conclusions

Motivated by current developments in the LNG industry, in this paper we develop a practical valuation model of LNG operations with a focus on the valuation of downstream storage. The unique aspect of our model is the integration of a closed-queuing-network model of natural-gas liquefaction and LNG shipping with an inventory-release model of LNG regasification and natural-gas sale into the wholesale spot market. We characterize the structure of the optimal

inventory-management policy at the downstream facility, and use our model to quantify the value and benefit of downstream storage for different players. Our results are relevant both at the theoretical and practical levels. Our theoretical analysis of the structure of the inventory-management policy extends in a significant manner the work of Secomandi (2007) to account for the interaction of shipping logistics and the operational parameters of the regasification facility. The resulting basestock-target structure is nontrivial, because when it is optimal to sell, it is not necessarily optimal to sell the entire available inventory. While our focus has been on LNG operations, this analysis remains relevant for other storable-commodity industries that exhibit uncertainty in the commodity production and distribution process, e.g., random yield, and/or where storage can be used to take advantage of spot-price fluctuations. On a practical level, our model can be used by downstream terminal operators and LNG players to assess the value of leasing contracts on regasification facilities. Our quantification of the value and benefit of storage also brings to light the importance of developing tactical-control software to optimize the management of LNG inventoried at regasification terminals.

We also believe that our model can be used in practice as an economic valuation model by different parties involved in the development of LNG projects. We explain this statement below. In the U.S., LNG is priced off the wholesale natural-gas price at a specific location. The U.S. LNG profitability-threshold is between \$2.70/MMBtu and \$3.30/MMBtu (Kaminski and Prevatt 2004), and the NYMEX futures curve as of 3/26/2007 fluctuates between \$6.85/MMBtu and \$9.59/MMBtu. Hence delivering LNG to the eastern U.S. promises hefty margins, irrespective of the specific operational configuration of a given project. Nevertheless, developing an LNG project is a long, complex, and expensive process (Greenwald 1998, Chapters 4 and 5). The operational length of a typical LNG project, once the necessary capacity has been installed, is between twenty and thirty years. In addition, five to ten years can be spent in negotiating the details of the project configuration and building the needed capacity. Several parties are involved in this process: State agencies of the producing country, the producer and the buyer, ship builders, several contractors, banks, and financiers. Numerous contracts are set up to govern the relationships among these parties: concession/production-sharing/development arrangements, operating/joint-venture agreements, the front end engineering and design contract, engineering, procurement and construction contracts, and the sale and purchase agreement (SPA). Negotiation of these contracts is complex and time consuming, and involves the building of intricate business relationships among the several parties involved. It is estimated that \$200 million or more can easily be spent in arriving at the final go/no-go decision point. The implication of these considerations is well-stated by Flower (1998, p. 120, emphasis added):

> [An] *economic model* of the project's expected *cash flow* is an important tool to have available as the project is developed. It will allow the economic returns of the project to be monitored [...] as negotiations progress with the LNG buyers and as arrangements with the host government and lenders are put in place. [...] [A] common project model [...] can be used *jointly* by the sponsors as a tool to aid decision-making during its development. For example, a joint model can be used to evaluate proposals from the buyers during negotiation of the SPA. This can help frame responses on issues such as volumes, build-up and price.

In other words, while it may be possible to make a fairly good decision on whether a particular project should be a go or a no-go based on simple calculations, the process of project development and contract negotiation can benefit from the availability of a more detailed mathematical model of the project operations. It is in this sense that we believe that our model can be useful in practice, in addition to the valuation of downstream storage: (i) it generates the physical and cash flows corresponding to a given LNG project, and (ii) evaluates them consistently with the market information expressed by traded natural-gas futures and options prices.

One could apply our model to study other interesting developments that are currently occurring in the LNG industry. For example, the number of new LNG ships on order is growing and the composition of the world LNG fleet is changing. Cho et al. (2005) report that while the most common size for LNG ships is between 135,000 to 145,000CM, "Qatargas II has [...] set [a] milestone by ordering eight large [LNG carriers]: four of 209,000[CM] and four of 216,000[CM]." They also state that these vessels "will exceed the largest LNG ship under construction at Chantiers de l'Atlantique, a 153,000 [CM] vessel ordered by Gas de France (GdF)." These larger vessel sizes reflect the explicit choice of some industry players to pursue and exploit shipping economies of scale. While these are well understood in practice, their throughput effects in terms of losses relative to installed capacity are not usually discussed in the industry press, e.g., in Cho et al. (2005). Here, it would be interesting to study the net benefit of shipping economies of scale, since employing larger ships may yield significant reductions in throughput, and, hence, system value, due to the well-known effects of increased transfer batches in manufacturing environments (see, e.g., Cachon and Terwiesch 2006, p. 92). Our model can be used to study these aspects.

Companies have also developed new regasification options, most notably Excelerate Energy Energy-Bridge technology, which allows specially designed tankers to regasify LNG on-board and off-load it to the existing pipeline grid by underwater pipelines (Jensen 2003, Gold 2004; see also the website of Excelerate Energy, www.excelerateenergy.com). This technology is now operational at the Gulf of Mexico Energy-Bridge Deepwater Port, located

100 miles off the Louisiana Cost, and a new development project has been approved in Massachusetts. While this technology removes the need for costly, and sometimes publicly opposed, on-land regasification terminals, it requires longer regasification time aboard the tankers, and increased tanker costs. Hence, industry players face the challenge of assessing its value relative to that of conventional terminal-based regasification technology. Our model can be modified to perform this comparison.

Our model can be extended and improved in other directions as well. For example, our shipping model can accommodate different vessel types and operating conditions, e.g., vessel speeds. Hence, one could extend our inventory-release model to allow tactical fleet control in response to market or inventory conditions. Also, while the strength of our approach is its practicality, our integration of the shipping and inventory-release models reflects a modeling simplification. Removing this approximation would require modeling the LNG operations as a stochastic dynamic program that at each point in time also keeps track of the state of the shipping system, the status of the liquefaction facility, and the inventory available at the upstream location. Such a model would be difficult to solve optimally using standard techniques, and its formulation and solution is a challenging area for additional research. Finally, LNG chains typically encompass, among other parties, a seller and a buyer, whose commercial relationships, as discussed above, are governed by an SPA. In this paper we have considered the benefit of downstream storage for different players based on exogenously specified LNG price functions. More broadly, the study of contractual issues related to the interactions between the LNG seller and buyer is a promising area for further research.

3.6 Appendix of Chapter 3

Proof of Proposition 5 (Optimal value-function). (i) (Concavity) By induction. Given the convexity assumption on function $h_{T+1}(\cdot)$, the property clearly holds in stage $T+1$. Make the induction hypothesis that the property also holds in stages $t+1,\dots,T$. Hence, $V_{t+1}(x_{t+1},p_{t+1})$ is concave in x_{t+1} for given p_{t+1}. This implies that $E_t^*[V_{t+1}(x_{t+1},\tilde{p}_{t+1})]$ is concave in x_{t+1} given p_t. Consider stage t and fix $p_t\in\mathcal{P}_t$, $u_t\in\mathcal{U}_t$, x_t^1, $x_t^2\in\mathcal{X}$ with $x_t^1\neq x_t^2$, and q_t^1, $q_t^2\in\mathcal{Q}(x_t,u_t)$ with $q_t^1\neq q_t^2$. Define $x_{t+1}^1:=x_t^1+u_t-q_t^1$, $x_{t+1}^2:=x_t^2+u_t-q_t^2$, and $x_{t+1}^\theta:=\theta x_{t+1}^1+(1-\theta)x_{t+1}^2$ for some $\theta\in[0,1]$. Since x_{t+1}^1, $x_{t+1}^2\in\mathcal{X}$, convexity of $\mathcal{X}$ implies that $x_{t+1}^\theta\in\mathcal{X}$. The concavity of $E_t^*[V_{t+1}(x_{t+1},\tilde{p}_{t+1})]$ in x_{t+1} for given p_t and the linearity of expectation imply that

$$E_t^*[V_{t+1}(x_{t+1}^\theta,\tilde{p}_{t+1})]\geqslant\theta E_t^*[V_{t+1}(x_{t+1}^1,\tilde{p}_{t+1})]+(1-\theta)E_t^*[V_{t+1}(x_{t+1}^2,\tilde{p}_{t+1})]. \tag{3.4}$$

Thus, the definitions of x_{t+1}^1 and x_{t+1}^2 imply that $E_t[V_{t+1}(x_t+u_t-q_t,\tilde{p}_{t+1})]$ is jointly concave in x_t and q_t for given p_t and u_t, and so is $\nu_t(x_t,q_t,p_t,u_t)$. This property, the convexity of set $\mathcal{A}(u_t)$, and Proposition B-4 in Heyman and Sobel (2004, p. 525) imply that $v_t(x_t,p_t,u_t) = \max_{q_t\in\mathcal{Q}(x_t,u_t)}\nu_t(x_t,q_t,p_t,u_t)$ is concave in x_t for given p_t and u_t. Concavity of $V_t(x_t,p_t)$ in x_t for given p_t follows since $V_t(x_t,p_t)=E[v_t(x_t,p_t,\tilde{u}_t)]$, and the property holds in all stages by the principle of mathematical induction.

(ii) (Supermodularity). By induction. Consider stage $T+1$. Since $p_{T+1}x_{T+1}$ is supermodular in $(x_{T+1},p_{T+1})\in\Re^2$ (see Topkis 1998, Example 2.6.2(c), p. 42), by the assumption that $g_{T+1}(p_{T+1})$ increases in p_{T+1}, it follows easily that $g_{T+1}(p_{T+1})x_{T+1}$ is supermodular in (x_{T+1},p_{T+1}) on the sublattice $\mathcal{X}\times\mathcal{P}_{T+1}\subset\Re^2$. Then, by Lemma 2.6.1(a)-(b) in Topkis (1998, p. 49), $V_{T+1}(x_{T+1},p_{T+1})$ is also supermodular in $(x_{T+1},p_{T+1})\in\mathcal{X}\times\mathcal{P}_{T+1}$ because $1-\varphi^R>0$ and $h_{T+1}(x_{T+1})$ is trivially supermodular in $(x_{T+1},p_{T+1})\in\mathcal{X}\times\mathcal{P}_{T+1}$. Make the induction hypothesis that the property also holds in stages $t+1,\ldots,T$. Consider stage t and fix $u_t\in\mathcal{U}_t$. Given $x_t\in\mathcal{X}$ and $q_t\in\mathcal{Q}(x_t,u_t)$, define $\phi_{t+1}(x_t,q_t,p_{t+1},u_t):=V_{t+1}(x_t+u_t-q_t,p_{t+1})$. Note that $\mathcal{A}(u_t)\subset\Re^2$ is a sublattice of $\Re^2$.

The concavity of $V_{t+1}(x_{t+1},p_{t+1})$ in $x_{t+1}\in\mathcal{X}$ for given $p_{t+1}\in\mathcal{P}_{t+1}$, established in part (i) of this proposition, and Lemma 2.6.2(b) in Topkis (1998, p. 50) imply that $\phi_{t+1}(x_t,q_t,p_{t+1},u_t)$ is supermodular in (x_t,q_t) on $\mathcal{A}(u_t)$ for given p_{t+1} and u_t. Denote $z_t=(x_t,q_t)$ and pick arbitrary z_t' and z_t'' both in $\mathcal{A}(u_t)$. Supermodularity of $\phi_{t+1}(z_t,p_{t+1},u_t)$ in $z_t\in\mathcal{A}(u_t)$ for given p_{t+1} and u_t means that

$$\phi_{t+1}(z_t',p_{t+1},u_t)-\phi(z_t'\wedge z_t'',p_{t+1},u_t)\leqslant\phi_{t+1}(z_t'\vee z_t'',p_{t+1},u_t)-\phi_{t+1}(z_t'',p_{t+1},u_t), \tag{3.5}$$

where $z_t'\wedge z_t''\equiv(\min\{x_t',x_t''\},\min\{q_t',q_t''\})$ and $z_t'\vee z_t''\equiv(\max\{x_t',x_t''\},\max\{q_t',q_t''\})$.

For $p_t\in\mathcal{P}_t$, define $\psi_t(z_t,p_t,u_t):=E[\phi_{t+1}(z_t,\tilde{p}_{t+1},u_t)|\tilde{p}_t=p_t]$. Pick p_t', $p_t''\in\mathcal{P}_t$ and, without loss of generality, assume that $p_t'\leqslant p_t''$. It will now be established that $\psi_t(z_t,p_t,u_t)$ is supermodular in $(z_t,p_t)\in\mathcal{A}(u_t)\times\mathcal{P}_t$, a sublattice of $\Re^3$, for given u_t, which means that the following inequality holds:

$$\psi_t(z_t',p_t',u_t)+\psi_t(z_t'',p_t'',u_t)\leqslant\psi_t(z_t'\vee z_t'',p_t'\vee p_t''\equiv p_t'',u_t)+\psi_t(z_t'\wedge z_t'',p_t'\wedge p_t''\equiv p_t',u_t). \tag{3.6}$$

Either (1) z_t' and z_t'' are ordered or (2) they are unordered.

Consider case (1). Assume that $z_t'\leqslant z_t''$, which means that $z_t'\wedge z_t''=z_t'$ and $z_t'\vee z_t''=z_t''$. Hence, it holds that $\phi_{t+1}(z_t',p_{t+1},u_t)=\phi_{t+1}(z_t'\wedge z_t'',p_{t+1},u_t)$ and $\phi(z_t'\vee z_t'',p_{t+1},u_t)=\phi(z_t'',p_{t+1},u_t)$, $\forall p_{t+1}\in\mathcal{P}_{t+1}$. It follows that (3.5) holds as an equality and so does (3.6)

because

$$\begin{aligned}
& \psi_t(z'_t, p'_t, u_t) - \psi_t(z'_t \wedge z''_t, p'_t, u_t) \\
= \; & E[\phi_{t+1}(z'_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z'_t \wedge z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p'_t] && (3.7) \\
= \; & E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p'_t] && (3.8) \\
= \; & E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p''_t] && (3.9) \\
= \; & \psi_t(z'_t \vee z''_t, p''_t, u_t) - \psi_t(z''_t, p''_t, u_t). && (3.10)
\end{aligned}$$

Consider case (2). Assume that $x'_t > x''_t$ and $q'_t < q''_t$. Thus, it holds that $(x'_t \vee x''_t) - (q'_t \vee q''_t) = x'_t - q''_t > x''_t - q''_t$. Since $V_{t+1}(x_{t+1}, p_{t+1})$ is supermodular in (x_{t+1}, p_{t+1}) it has increasing differences in p_{t+1} by Theorem 2.6.1 in Topkis (1998), i.e., given x^1_{t+1}, $x^2_{t+1} \in \mathcal{X}$, with $x^1_{t+1} < x^2_{t+1}$, the quantity $V_{t+1}(x^2_{t+1}, p_{t+1}) - V_{t+1}(x^1_{t+1}, p_{t+1})$ increases in p_{t+1}. By letting $x^2_{t+1} := x'_t + u_t - q''_t$ and $x^1_{t+1} := x''_t + u_t - q''_t$, this is equivalent to stating that $\phi_{t+1}(z'_t \vee z''_t, p_{t+1}, u_t) - \phi_{t+1}(z''_t, p_{t+1}, u_t)$ increases in p_{t+1} for given u_t. This statement, the assumption that the distribution of $\tilde{p}_{t+1}|p_t$ stochastically increases in p_t, and Corollary 3.9.1 in Topkis (1998, p. 161) imply

$$\begin{aligned}
& E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p'_t] \\
\leqslant \; & E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p''_t]. && (3.11)
\end{aligned}$$

Expression (3.6) is true because proceeding similarly to the proof of Theorem 3.10.1 in Topkis (1998) yields that

$$\begin{aligned}
& \psi_t(z'_t, p'_t, u_t) - \psi_t(z'_t \wedge z''_t, p'_t, u_t) \\
= \; & E[\phi_{t+1}(z'_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z'_t \wedge z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p'_t] && (3.12) \\
\leqslant \; & E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p'_t]; \text{ by (3.5)} \\
\leqslant \; & E[\phi_{t+1}(z'_t \vee z''_t, \tilde{p}_{t+1}, u_t) - \phi_{t+1}(z''_t, \tilde{p}_{t+1}, u_t) | \tilde{p}_t = p''_t]; \text{ by (3.11)} \\
= \; & \psi_t(z'_t \vee z''_t, p''_t, u_t) - \psi_t(z''_t, p''_t, u_t). && (3.13)
\end{aligned}$$

Hence, $\psi_t(z_t, p_t, u_t)$ is supermodular in $(z_t, p_t) \in \mathcal{A}(u_t) \times \mathcal{P}_t$ for given u_t, and $\delta_t E[V_{t+1}(x_t + u_t - q_t, \tilde{p}_{t+1}) | \tilde{p}_t = p_t]$ is supermodular in $(x_t, q_t, p_t) \in \mathcal{A}(u_t) \times \mathcal{P}_t$ for given u_t by Lemma 2.6.1(a) in Topkis (1998, p. 49). The term $g_t(p_t)(1 - \varphi^R)q_t$ is supermodular in (q_t, p_t) on $\mathcal{Q}(x_t, u_t) \times \mathcal{P}_t$, a sublattice of $\Re^2$, since $g_t(p_t)$ increases in p_t (by assumption) and $1 - \varphi^R$ is positive. Since $h_t(x_t)$ is trivially supermodular in $(x_t, q_t, p_t) \in \mathcal{A}(u_t) \times \mathcal{P}_t$ for given u_t, it follows that $\nu_t(x_t, q_t, p_t, u_t)$ is supermodular in $(x_t, q_t, p_t) \in \mathcal{A}(u_t) \times \mathcal{P}_t$, a sublattice of $\Re^3$, for given u_t. Theorem 2.7.6 in Topkis (1998, p. 70) then implies that $v_t(x_t, p_t, u_t)$ is supermodular in (x_t, p_t) on $\mathcal{X} \times \mathcal{P}_t$ for given u_t. Since this property holds for each $u_t \in \mathcal{U}_t$,

Corollary 2.6.2 in Topkis (1998, p. 50) yields that $V_t(x_t, p_t)$ is supermodular in (x_t, p_t) on $\mathcal{X} \times \mathcal{P}_t$. Thus, this property holds in all stages by the principle of mathematical induction. $\square$.

Proof of Proposition 6 (Basestock target). In arbitrary stage $t \in \mathcal{T}$, fix $p_t \in \mathcal{P}_t$, pick any two distinct values $y_t^1,\ y_t^2 \in \mathcal{X}$ such that $y_t^1 < y_t^2$, and define $\Delta(y_t^1, y_t^2) := y_1^2 - y_t^1$, so that $y_t^2 = y_t^1 + \Delta(y_t^1, y_t^2)$. It holds that

$$\begin{aligned}
\check{\nu}_t(y_t^1, q_t, p_t) &= g_t(p_t)(1-\varphi^R)q_t + \delta_t E_t^*[V_{t+1}(y_t^1 - q_t, \tilde{p}_{t+1})] \\
&= g_t(p_t)(1-\varphi^R)[q_t + \Delta(y_t^1, y_t^2)] + \delta_t E_t^*[V_{t+1}(y_t^2 - (q_t + \Delta(y_t^1, y_t^2)), \tilde{p}_{t+1})] \\
&\quad -g_t(p_t)(1-\varphi^R)\Delta(y_t^1, y_t^2) \\
&= \check{\nu}_t(y_t^2, q_t + \Delta(y_t^1, y_t^2), p_t) - g_t(p_t)(1-\varphi^R)\Delta(y_t^1, y_t^2). \qquad (3.14)
\end{aligned}$$

To interpret this equality, think of $\nu_t(y_t, q_t, p_t)$ as a function of q_t at different values of y_t for a given value of p_t. Equality (3.14) states that the graph of this function at y_t^1 is the graph of this function at y_t^2 shifted to the left by distance $\Delta(y_t^1, y_t^2)$, and down by distance $g_t(p_t)(1-\varphi^R)\Delta(y_t^1, y_t^2)$. Recall that we define $\check{q}_t^\diamond(y_t, p_t)$ as the largest quantity in set $\arg\max_{q_t \in \check{\mathcal{Q}}_t(y_t)} \check{\nu}_t(y_t, q_t, p_t)$. Define $\underline{y}_t$ as the largest inventory level $y_t \in \mathcal{X}$ such that $\check{q}_t^\diamond(y_t, p_t) = 0$. Note that $\underline{y}_t$ exists since $\check{q}_t^\diamond(0, p_t) = 0$. If $\underline{y}_t = \overline{x}$ then $b_t(p_t) = \overline{x}$. If $\underline{y}_t \neq \overline{x}$, consider any two distinct inventory levels $y_t^1,\ y_t^2 \in (\overline{y}_t, \overline{x}]$ such that $y_t^2 > y_t^1 > \underline{y}_t$. Consider a value $q_t \in \check{Q}(y_t^1) \equiv [0, y_t^1]$ and notice that $q_t + \Delta(y_t^1, y_t^2)$ is feasible at y_t^2 since $q_t + \Delta(y_t^1, y_t^2) \in [\Delta(y_t^1, y_t^2), y_t^1 + \Delta(y_t^1, y_t^2)] = [\Delta(y_t^1, y_t^2), y_t^2] \subset [0, y_t^2] \equiv \check{\mathcal{Q}}_t(y_t^2)$. Since $\check{q}^\diamond(y_t^1, p_t) > 0$ and $\check{\nu}_t(y_t, q_t, p_t)$ is concave in q_t, equality (3.14) implies that $\check{q}^\diamond(y_t^2, p_t) \in [\Delta(y_t^1, y_t^2), y^2]$. In fact, this equality implies that $\check{q}^\diamond(y_t^2, p_t) = \check{q}^\diamond(y_t^1, p_t) + \Delta(y_t^1, y_t^2)$. Therefore

$$y_t^2 - \check{q}^\diamond(y_t^2, p_t) = y_t^2 - \Delta(y_t^1, y_t^2) - \check{q}^\diamond(y_t^1, p_t) = y_t^1 - \check{q}^\diamond(y_t^1, p_t), \qquad (3.15)$$

so that releasing amounts $\check{q}^\diamond(y_t^2, p_t)$ and $\check{q}^\diamond(y_t^1, p_t)$ at y_t^2 and y_t^1, respectively, leads to the same inventory level. Concavity of $\check{\nu}_t(y_t, q_t, p_t)$ in q_t, equality (3.14), and the definition of $\underline{y}_t$ imply that

$$\check{q}^\diamond(y_t^1, p_t) \leqslant \Delta(\underline{y}_t, y_1), \qquad (3.16)$$

because otherwise $\check{q}_t^\diamond(\underline{y}_t, p_t) \neq 0$, a contradiction. Thus, it holds that

$$y_t^1 - \check{q}_t^\diamond(y_t^1, p_t) \geqslant y_t^1 - \Delta(\underline{y}_t, y_t^1) = \underline{y}_t,$$

and therefore $b_t(p_t) \geqslant \underline{y}_t$. We now show that in fact $b_t(p_t) = \underline{y}_t$. Select y_t^1 to be arbitrarily close to $\underline{y}_t$, i.e., $y_t^1 = \underline{y}_t + \epsilon$ with $\epsilon \in \Re_+$ and arbitrarily small. From inequality (3.16) and $\epsilon \equiv \Delta(\underline{y}_t, y_t^1)$, we obtain that

$$\lim_{\epsilon \to 0} \check{q}_t^\diamond(\underline{y} + \epsilon, p_t) \leqslant \lim_{\epsilon \to 0} \epsilon = 0 \Rightarrow \lim_{\epsilon \to 0} \check{q}_t^\diamond(\underline{y} + \epsilon, p_t) = 0 \text{ since } \check{q}_t^\diamond(\underline{y} + \epsilon, p_t) \geqslant 0.$$

Thus, it also holds that $\lim_{\epsilon \to 0} \underline{y}_t + \epsilon - \breve{q}_t^{\diamond}(\underline{y} + \epsilon, p_t) = \underline{y}_t$, and by equation (3.15) $y_t^2 - \breve{q}_t^{\diamond}(y_t^2, p_t) = \underline{y}_t$ for all $y_t^2 \in (\underline{y}_t + \epsilon, \overline{x}]$, which implies that $b_t(p_t) = \underline{y}_t$. $\square$

Computation of the shipping operating cost in §3.4. Cho et al. (2005) report a charter rate of \$65,000/day for a 145,000CM vessel covering a one-way distance of 7,000NM at a speed of 19 knots during a period of twenty years (see Table 3 of their paper for specific details). These authors also state that the charter rate is about 68% of the total shipping costs, with the remaining 32% being the shipping operating cost. This implies that the shipping operating cost per day for this vessel under the stated conditions is $\$65,000 \cdot 0.32/0.68 = \$30,588.235$. Assuming that the ship is operated 365 days per year, this translates to \$11,164,706 per year. The operating conditions on which these computations are based are identical to those used in our numerical study.

Explanation of the decreasing value of storage for "high" throughput in §3.4. Consider a degenerate unloading random variable equal to u. Momentarily, impose the constraint that the inventory level in period $T+1$ be zero, i.e., $x_{T+1} = 0$. Suppose that $u = \overline{q}$, so that the rate into the terminal is equal to the maximum rate out of it. In this case, the value of storage must be zero because no amount of LNG can be stored in any period. Thus, when u is sufficiently high, as it approaches $\overline{q}$, from below since $u \leqslant \overline{q}$, the value of storage decreases to zero. Now, remove the constraint $x_{T+1} = 0$, so that the only conditions imposed on x_{T+1} are $0 \leqslant x_{T+1} \leqslant \overline{x}$. Make the realistic assumption that $\overline{q}T \geqslant \overline{x}$, i.e., a full terminal in period 1 can be emptied by time $T+1$. If $u = \overline{q}$, any amount of LNG not released in some period $t \in \mathcal{T}$ must be stored until period $T+1$, and the maximum amount of stored LNG during the *entire* planning horizon is $\min(uT, \overline{x}) = \min(\overline{q}T, \overline{x}) = \overline{x}$. Thus, the value of storage for any level of throughput that allows storing at least an amount $\overline{x}$ of LNG during the entire planning horizon, i.e., for any $u \leqslant \overline{q}$ such that $uT \geqslant \overline{x}$, must be at least the one obtainable when $u = \overline{q}$. In other words, $u = \overline{q}$ is the level of throughput that *minimizes* the value of storage among all those that satisfy $uT \geqslant \overline{x}$. Therefore, as $u \geqslant \overline{x}/T$ approaches $\overline{q}$, obviously from below, the value of storage decreases in a neighborhood of $\overline{q}$. When the unloading random variable is not degenerate, explaining the decreasing value of storage after some level of throughput is more involved, but the main intuition provided here remains relevant.

Chapter 4

The Importance of Modeling Shipping Variability When Valuing Downstream LNG Storage

4.1 Introduction

In Chapter 3, we use a closed-queuing-network (CQN) model to represent the upstream LNG production and ocean shipping logistics to a downstream regasification facility in the U.S., which sells to the wholesale spot market. To model supply variability, the shipping model has two steps: the first step is to analytically compute the stationary distribution $\pi(n)$ for the ship locations within the network at a given instant in time; and the second step is to generate the probability distribution of the unloaded amount to the downstream LNG storage facility during a certain time period (e.g. a month) using the rolling-forward method (see section 3.2.1 for details). We thus assume that in every period the unloading process is independently and identically distributed (i.i.d.). Then we feed this i.i.d. unloaded amount distribution into the inventory release model together with the stochastic price from the price trinomial lattice model to value the option to store.

In this chapter, we will focus on studying how important it is to model shipping variability while valuing or managing a downstream LNG storage facility. Table 3.3 in Chapter 3 shows that if we ignore the loading/unloading congestion and shipping time uncertainty, we will overvalue the system throughput, especially when there are several LNG projects sharing the same loading/unloading ports, which is always the case in the LNG industry. (For example, Qatar supplies LNG to Asia, Europe and North America.) At the downstream side, U.S. imports LNG from Trinidad and Tobago, Egypt, Qatar among other countries, and several LNG projects share the same unloading port at Lake Charles, Louisiana. Therefore Table

3.3 shows that it is important to model the shipping variability. To validate the shipping model presented in Chapter 3, we still need to answer the following two questions.

Did we over-simplify the shipping model? In general, if there are more ships arriving in the current period, we will expect fewer ships arriving in the next period; there is dependence between the unloaded amounts in consecutive periods. One may thus question the i.i.d. assumption. To answer this question, we compare the storage values using the i.i.d. unloaded amount distribution with values found using a simulation model of the shipping process. If the two values are very close, we can conclude that the shipping model presented in Chapter 3 is a good representation of the realistic shipping process to value LNG storage facilities.

The second question is: can we further simplify the shipping model, if it is not already over-simplified? Could we simply use the throughput in the inventory release model instead of the unloaded amount distribution? We only need the first step of the shipping model, $\pi(n)$, to compute the throughput; it is much easier and much faster. To answer this question, we compare the storage values using the constant throughput with the ones using the unloaded amount distribution or simulation. If the storage values are very close, then we can greatly simplify the shipping model.

We present our models in section 4.2, and empirically assess the performance of our models in section 4.3 to answer the above questions.

4.2 Models

We consider the same closed-queuing-network (CQN) presented in Figure 3.2, and model the loading and unloading blocks as first-come-first-served (FCFS) exponential queues, and the transit blocks as ample-server (AS) exponential or multi-stage Coxian queues. To test whether the shipping model in Chapter 3 is over-simplified or can be further simplified, we propose a benchmark model as follows.

Benchmark Model. Instead of using analytical formulas to compute the unloaded amount distribution, we use simulation to generate sample paths of unloaded cargos. When the decision horizon is 20 years with monthly intervals, our shipping simulation model will output 10,000 (sample size) sample paths each consisting of 240 numbers, which are the unloaded cargos from period 1 to period 240. Then for each sample path we use the inventory release model to determine the optimal sale quantity in every period, given the price realization and inventory level, and given *perfect* information of future LNG cargo arrivals. Therefore the Benchmark Model provides an *upper bound* of the LNG storage value.

Next, we define the following models and compare them with the Benchmark Model.

Model 1: We use the rolling-forward method to generate the unloaded amount distribu-

tion, and use this distribution to value the option to store. This is the model we present in Chapter 3. We compare this model with the Benchmark Model to test if Model 1 is effective for valuation purposes.

Model 2 has two steps. First, we input the unloaded amount distribution found by the rolling-forward method to the inventory release model to obtain the base-stock levels given each price realization and available inventory level at each period. Thus, Model 2 is a simulation-based verification of the dynamic program value found by Model 1.

Model 3: We use *constant* shipping throughput in the inventory release model for valuation purpose. We use the CQN shipping model to compute the throughput taking the port congestion into account. After we get the throughput, we assume that a fixed number of cargos arrive at the downstream LNG storage facility in each period, which is equal to the throughput, u; in the inventory release model, the arriving supply rate u has no uncertainty. However, the spot price is still stochastic. We use the trinomial lattice to model the price uncertainty. Similar to Model 1, the stochastic-dynamic-programming valuation model is defined as

$$\begin{aligned} V_{T+1}(x_{T+1}, p_{T+1}) &:= g_{T+1}(p_{T+1})(1-\varphi^R)x_{T+1} - h_{T+1}(x_{T+1}), \forall x_{T+1} \in \mathcal{X},\ p_{T+1} \in \mathcal{P}_{T+1} \\ V_t(x_t, p_t) &= \max_{q_t \in \mathcal{Q}(x_t,u)} \nu_t(x_t, q_t, p_t, u) \\ \nu_t(x_t, q_t, p_t, u) &:= g_t(p_t)(1-\varphi^R)q_t - h_t(x_t) - c_t^U u + \delta_t E_t^*[V_{t+1}(x_t + u - q_t, \tilde{p}_{t+1})]. \end{aligned}$$

Model 4: We use Model 3 to get the base-stock levels for each price realization in each period. Then we simulate the resulting inventory control policy, as in Model 2.

4.3 Numerical Examples

We now use numerical exmples to compare Model 1-4 with the Benchmark Model. We use the same experiment setting as in Chapter 3, summarized in Table 3.2 (see section 3.4 operational parameters and operating costs part for details).

First, we check the throughputs (mean LNG supply rates) of these models. Model 1 and 3 use the unloaded amount distribution, and constant throughput, respectively, to value the option to store. We expect the throughputs of these two models to be exactly the same; both of them are the results of same analytical formulas. Similary, the Benchmark Model, Model 2, and Model 4 use the simulated shipping process, so we expect the throughputs of these three models to be the same since we use common random variables. Column 3 of Table 4.1 shows that the throughputs using analytical formulas and simulation are very close. This validates the accuracy of the shipping models.

System Value. Table 4.1 reports the system values (the discounted expected cash flows) of the five models over twenty years when the price volatility is set to 0.7, and Table 4.2 displays the system values of Model 1-4 as fractions of those of the Benchmark Model. We can see that the system values using different models are very close (less than 1% difference) given any system configuration (fleet size and storage size). This means that we can use any of these models to estimate the system value.

Storage Value is defined as the system value with the storage facility less the system value without storage. Generally, the system values are several billion dollars, and storage values are several hundred million dollars. Table 4.3 shows the storage values (unit: million dollars), and Table 4.4 displays the storage values of Model 1-4 as fractions of those of benchmark model, when the price volatility is set to 0.7.

The average value of Model 1 is higher than that of Model 3 (0.9894 vs. 0.9527), and very close to those of Models 2 and 4. Except for the three special cases, fleet size = 1 ship and storage size = 3, 4, and 5 cargos, the storage value of Model 3 is significantly lower than those of Model 1 and the Benchmark Model. This result shows that generally Model 1 using the unloaded amount distribution is better for valuation purposes than Model 3, since the storage value estimated by Model 1 is closer to both the values of the Benchmark Model (0.9894 vs. 1) and those of Model 2 (0.9894 vs. 0.9850).

"Average" in Column 1 of Table 4.4 is defined as the sum of all storage values using Model 1-4 divided by the sum of those of the Benchmark Model. The average value of Model 2 is slightly higher than that of Model 4 (0.9850 vs. 0.9811), and they are very close to the upper bound (less than 2% difference). After further comparing the storage values for the different system configurations (fleet size and storage size), we find that the storage values of Model 4 are much lower than those of Model 2 (0.949 vs. 0.969; 0.921 vs. 0.965, 0.963 vs. 0.975), when the fleet size is 1 ship and the storage size is 4 and 5 cargos; and when the fleet size is 2 ships and the storage size is 5 cargos. Except for these special cases, the storage values of these two models are very close. In these special cases, the ratio of shipping supply rate to storage size is low. This means that most of the time the maximum send-out rate and storage space are not binding constraints, and the storage managers have more freedom to play the store-low-sell-high game: optional sales play a more important role. Therefore in these cases it is better to consider more detailed available inventory level when generating the inventory control policy using the unloaded amount distribution rather than using the constant throughput.

Intrinsic Value is defined as the storage value when the prices are deterministic; we simply set the price volatility to zero. Table 4.5 shows the intrinsic values (unit: million dollars), and Table 4.6 displays the intrinsic value of Model 1-4 as fractions of those of the

Table 4.1: Comparison of the system values of the Benchmark Model and Models 1-4 (unit: billion dollars)

Model	# of Ships	Throughput (MTPA)	Storage Size (# of Cargos)					
			0	1	2	3	4	5
Benchmark Model	1	0.7424	3.765	3.870	3.968	4.056	4.134	4.202
	2	1.4847	7.527	7.634	7.739	7.839	7.933	8.021
	3	2.2190	11.247	11.355	11.461	11.566	11.667	11.763
	4	2.9574	14.993	15.101	15.208	15.314	15.418	15.519
	5	3.6893	18.697	18.804	18.912	19.019	19.124	19.227
	6	4.4092	22.348	22.456	22.563	22.669	22.774	22.876
Model 1	1	0.7412	3.757	3.861	3.956	4.041	4.116	4.181
	2	1.4814	7.509	7.615	7.719	7.817	7.909	7.995
	3	2.2195	11.250	11.357	11.463	11.566	11.665	11.759
	4	2.9541	14.971	15.079	15.186	15.291	15.394	15.493
	5	3.6831	18.664	18.771	18.879	18.985	19.089	19.189
	6	4.3944	22.264	22.371	22.478	22.583	22.686	22.785
Model 2	1	0.7424	3.765	3.869	3.963	4.048	4.123	4.187
	2	1.4847	7.527	7.633	7.736	7.832	7.923	8.009
	3	2.2190	11.247	11.355	11.460	11.562	11.660	11.753
	4	2.9574	14.993	15.100	15.207	15.313	15.414	15.512
	5	3.6893	18.697	18.804	18.911	19.018	19.121	19.220
	6	4.4092	22.348	22.455	22.562	22.667	22.768	22.867
Model 3	1	0.7412	3.757	3.858	3.953	4.042	4.125	4.202
	2	1.4814	7.509	7.609	7.710	7.806	7.901	7.990
	3	2.2195	11.250	11.350	11.451	11.552	11.647	11.742
	4	2.9541	14.971	15.072	15.172	15.273	15.373	15.468
	5	3.6831	18.664	18.764	18.864	18.965	19.065	19.165
	6	4.3944	22.264	22.363	22.463	22.563	22.663	22.762
Model 4	1	0.7424	3.765	3.869	3.964	4.047	4.115	4.167
	2	1.4847	7.527	7.633	7.736	7.833	7.922	8.003
	3	2.2190	11.247	11.355	11.460	11.562	11.660	11.752
	4	2.9574	14.993	15.100	15.207	15.312	15.414	15.512
	5	3.6893	18.697	18.804	18.912	19.018	19.121	19.219
	6	4.4092	22.348	22.455	22.562	22.667	22.768	22.866

Table 4.2: Comparison of Models 1-4 system values as fractions of those of the Benchmark Model

Models	# of Ships	Storage Size (# of Cargos)					
		0	1	2	3	4	5
Model 1	1	0.998	0.998	0.997	0.996	0.996	0.995
	2	0.998	0.998	0.997	0.997	0.997	0.997
	3	1.000	1.000	1.000	1.000	1.000	1.000
	4	0.999	0.999	0.999	0.999	0.998	0.998
	5	0.998	0.998	0.998	0.998	0.998	0.998
	6	0.996	0.996	0.996	0.996	0.996	0.996
Model 2	1	1.000	1.000	0.999	0.998	0.997	0.996
	2	1.000	1.000	1.000	0.999	0.999	0.998
	3	1.000	1.000	1.000	1.000	0.999	0.999
	4	1.000	1.000	1.000	1.000	1.000	1.000
	5	1.000	1.000	1.000	1.000	1.000	1.000
	6	1.000	1.000	1.000	1.000	1.000	1.000
Model 3	1	0.998	0.997	0.996	0.997	0.998	1.000
	2	0.998	0.997	0.996	0.996	0.996	0.996
	3	1.000	1.000	0.999	0.999	0.998	0.998
	4	0.999	0.998	0.998	0.997	0.997	0.997
	5	0.998	0.998	0.998	0.997	0.997	0.997
	6	0.996	0.996	0.996	0.995	0.995	0.995
Model 4	1	1.000	1.000	0.999	0.998	0.995	0.992
	2	1.000	1.000	1.000	0.999	0.999	0.998
	3	1.000	1.000	1.000	1.000	0.999	0.999
	4	1.000	1.000	1.000	1.000	1.000	1.000
	5	1.000	1.000	1.000	1.000	1.000	1.000
	6	1.000	1.000	1.000	1.000	1.000	1.000

Table 4.3: Comparison of the storage values of the Benchmark Model and Models 1-4 (unit: million dollars)

Models	# of Ships	Throughput (MTPA)	Storage Size (# of Cargos)				
			1	2	3	4	5
Benchmark Model	1	0.7424	105.4	203.6	291.4	369.3	437.2
	2	1.4847	107.1	211.9	312.1	406.2	494.4
	3	2.2190	107.5	214.2	318.6	419.4	515.8
	4	2.9574	107.6	214.9	321.2	425.3	526.4
	5	3.6893	107.6	215.1	322.0	427.3	530.1
	6	4.4092	107.5	214.8	321.2	425.9	527.9
Model 1	1	0.7412	104.0	199.4	284.4	359.3	423.9
	2	1.4814	106.8	210.2	308.1	399.9	486.4
	3	2.2195	107.4	213.5	316.4	415.1	509.4
	4	2.9541	107.6	214.7	320.1	422.7	521.9
	5	3.6831	107.6	214.9	321.2	425.0	525.5
	6	4.3944	107.5	214.3	319.5	422.1	521.3
Model 2	1	0.7424	103.9	198.5	282.7	357.7	422.0
	2	1.4847	106.6	209.4	305.6	396.0	481.9
	3	2.2190	107.4	213.1	315.2	412.4	505.5
	4	2.9574	107.6	214.5	319.6	421.5	519.0
	5	3.6893	107.6	214.9	320.9	424.1	523.1
	6	4.4092	107.4	214.0	318.9	420.4	518.6
Model 3	1	0.7412	100.9	196.3	285.5	368.3	445.5
	2	1.4814	100.8	201.7	297.1	392.4	481.6
	3	2.2195	100.7	201.5	302.2	397.5	492.8
	4	2.9541	100.6	201.1	301.7	402.2	497.4
	5	3.6831	100.3	200.6	301.0	401.3	501.6
	6	4.3944	99.8	199.6	299.4	399.2	499.0
Model 4	1	0.7424	103.7	199.0	281.7	350.3	402.5
	2	1.4847	106.7	209.2	306.5	395.3	476.1
	3	2.2190	107.4	213.1	314.8	412.7	504.4
	4	2.9574	107.6	214.5	319.6	420.9	518.9
	5	3.6893	107.6	214.9	320.9	424.1	522.5
	6	4.4092	107.1	213.9	318.9	420.4	518.4

Table 4.4: Comparison of the storage values of Models 1-4 as fractions of those of the Benchmark Model

Models	# of Ships	Storage Size (# of Cargos)				
		1	2	3	4	5
Model 1 (Average = 0.9894)	1	0.987	0.980	0.976	0.973	0.969
	2	0.996	0.992	0.987	0.984	0.984
	3	0.999	0.997	0.993	0.990	0.988
	4	1.000	0.999	0.997	0.994	0.991
	5	1.000	0.999	0.998	0.995	0.991
	6	0.999	0.998	0.995	0.991	0.987
Model 2 (Average = 0.9850)	1	0.986	0.975	0.970	0.969	0.965
	2	0.995	0.988	0.979	0.975	0.975
	3	0.999	0.995	0.989	0.983	0.980
	4	1.000	0.998	0.995	0.991	0.986
	5	1.000	0.999	0.997	0.992	0.987
	6	0.998	0.996	0.993	0.987	0.982
Model 3 (Average = 0.9527)	1	0.957	0.964	0.980	0.997	1.019
	2	0.941	0.952	0.952	0.966	0.974
	3	0.937	0.941	0.949	0.948	0.955
	4	0.934	0.936	0.939	0.946	0.945
	5	0.932	0.933	0.935	0.939	0.946
	6	0.928	0.929	0.932	0.937	0.945
Model 4 (Average = 0.9811)	1	0.983	0.977	0.967	0.949	0.921
	2	0.996	0.987	0.982	0.973	0.963
	3	0.999	0.995	0.988	0.984	0.978
	4	0.999	0.998	0.995	0.990	0.986
	5	1.000	0.999	0.997	0.992	0.986
	6	0.996	0.996	0.993	0.987	0.982

Table 4.5: Comparison of the intrinsic values of the Benchmark Model and Models 1-4 (unit: million dollars)

Models	# of Ships	Throughput (MTPA)	Storage Size (# of Cargos)				
			1	2	3	4	5
Benchmark Model	1	0.7424	79.8	157.4	231.7	300.3	358.5
	2	1.4847	80.3	159.8	237.9	314.3	388.2
	3	2.2190	80.4	160.5	239.9	318.2	395.4
	4	2.9574	80.4	160.8	240.8	320.1	398.5
	5	3.6893	80.5	160.8	241.0	320.8	399.7
	6	4.4092	80.4	160.7	240.8	320.3	398.7
Model 1	1	0.7412	79.1	156.4	229.2	294.2	345.2
	2	1.4814	80.1	159.1	237.1	313.3	386.7
	3	2.2195	80.4	160.3	239.2	317.3	394.3
	4	2.9541	80.4	160.7	240.4	319.3	397.5
	5	3.6831	80.5	160.8	240.9	320.3	398.9
	6	4.3944	80.4	160.8	240.7	320.0	397.9
Model 2	1	0.7424	78.9	155.4	228.4	294.3	345.2
	2	1.4847	80.1	159.0	235.8	310.9	382.1
	3	2.2190	80.4	160.2	239.1	317.0	393.6
	4	2.9574	80.4	160.6	240.3	319.1	397.0
	5	3.6893	80.5	160.8	240.8	320.2	398.7
	6	4.4092	80.4	160.7	240.7	319.9	398.0
Model 3	1	0.7412	75.4	148.8	220.8	291.1	360.3
	2	1.4814	75.4	150.7	224.1	297.5	369.4
	3	2.2195	75.3	150.6	225.8	299.2	372.5
	4	2.9541	75.2	150.3	225.5	300.7	373.9
	5	3.6831	75.0	150.0	225.0	300.0	375.0
	6	4.3944	74.6	149.3	223.9	298.5	373.2
Model 4	1	0.7424	78.5	154.3	227.8	293.2	342.2
	2	1.4847	80.1	158.3	235.2	310.2	383.5
	3	2.2190	80.4	160.2	238.1	315.2	391.3
	4	2.9574	80.4	160.6	240.3	317.9	395.2
	5	3.6893	80.5	160.8	240.8	320.1	397.1
	6	4.4092	80.4	160.7	240.7	319.9	398.0

Table 4.6: Comparison of the intrinsic values of Models 1-4 as fractions of those of the Benchmark Model

Models	# of Ships	Storage Size (# of Cargos)				
		1	2	3	4	5
Model 1 (Average = 0.9950)	1	0.992	0.994	0.989	0.979	0.963
	2	0.998	0.996	0.996	0.997	0.996
	3	1.000	0.999	0.997	0.997	0.997
	4	1.000	1.000	0.999	0.998	0.998
	5	1.000	1.000	0.999	0.999	0.998
	6	1.000	1.000	1.000	0.999	0.998
Model 2 (Average = 0.9932)	1	0.990	0.988	0.986	0.980	0.963
	2	0.998	0.995	0.991	0.989	0.984
	3	1.000	0.998	0.997	0.996	0.995
	4	1.000	0.999	0.998	0.997	0.996
	5	1.000	1.000	0.999	0.998	0.998
	6	1.000	1.000	1.000	0.999	0.998
Model 3 (Average = 0.9437)	1	0.945	0.946	0.953	0.969	1.005
	2	0.939	0.943	0.942	0.946	0.951
	3	0.936	0.938	0.941	0.940	0.942
	4	0.934	0.935	0.937	0.939	0.938
	5	0.932	0.933	0.933	0.935	0.938
	6	0.928	0.929	0.930	0.932	0.936
Model 4 (Average = 0.9908)	1	0.984	0.980	0.983	0.976	0.955
	2	0.998	0.991	0.989	0.987	0.988
	3	1.000	0.998	0.992	0.991	0.990
	4	1.000	0.999	0.998	0.993	0.992
	5	1.000	0.999	0.999	0.998	0.993
	6	0.999	1.000	1.000	0.999	0.998

Benchmark Model. Except for the special case, fleet size = 1 ship and storage size = 5 cargos, the intrinsic values using Model 1 are significantly higher than those of Model 3. The intrinsic values of Model 2 and 4 are very close (0.9950 vs. 0.9908), and very close to those of the Benchmark Model (less than 1% difference). This means that Model 1, using the unloaded amount distribution, is much better than Model 3, using the constant throughput for LNG storage valuation purposes, and both modeling methods of cargo arrivals can yield a near optimal inventory control policy when there is no price uncertainty (but price seasonality is still present).

4.4 Conclusions

In this chapter we examine how important it is to model the shipping variability while valuing and managing the downstream LNG storage; this is an extension of Chapter 3. Through a numerical study, we validate that the shipping model presented in Chapter 3 is a good representation of the shipping process, and further understand under what kind of circumstance the model of the shipping process can be simplified. The model using the unloaded amount distribution provides a good estimation of the storage value, and also yield a very good inventory control policy. The model using the constant throughput is not as good as the previous model for storage valuation purposes, but can be used to suggest the optimal inventory control policy, especially when the ratio of supply arrival rate to storage size is high.

Chapter 5

Conclusions

This thesis studies real option valuation problems of natural gas transport capacity and LNG downstream storage.

Natural gas pipeline capacity contracts are the essential assets for natural gas shippers. It is common for natural gas shippers to value contracts on this capacity by simple adaptations of financial spread option formulas, which do not fully account for the implications of the capacity limits and the network structure that distinguish these contracts. In contrast, we show that these operational features can be fully captured and integrated with financial considerations in a fairly easy and managerially significant manner by a model that combines linear programming and simulation (LPS). We also derive spread option and linear/convex programming based lower and upper bounds on the economic value of capacity that are easier to compute (at least approximately). Our lower bound (LB) model appears to be an enhanced version of models used in practice. Our upper bound (UB) model provides an additional valuation benchmark. Based on actual prices of traded natural gas futures and basis swaps, we show that an enhanced version of the common approach employed in practice can significantly undervalue natural gas pipeline network capacity relative to our LPS model. We derive pathwise estimators for the so called deltas and structurally characterize them. We also interpret them in a novel fashion as discounted expectations, under a specific weighing distribution, of the amounts of natural gas to be procured/marketed when optimally using pipeline capacity. Our model also exhibits promising financial (delta) hedging performance. Thus, our LPS model emerges as an easy to use and useful tool that natural gas shippers can employ to support their valuation and delta hedging decisions concerning natural gas pipeline network transport capacity contracts. Models currently used in practice (proxied by our LB model) can significantly undervalue network capacity when compared to our LPS model, by 5-25% in different cases. Our UB model considerably overvalues this capacity. The delta hedging performance of our LPS model is encouraging. Thus, our LPS model

emerges as a managerially relevant model both in terms of practical implementation and usage. Moreover, the insights that follow from our data analysis have broader significance and implications in terms of the management of real options beyond our specific application.

The U.S. Energy Information Administration forecasts that local natural-gas production will be soon unable to meet demand in the U.S. and most other developed countries, and, not coincidentally, expects liquefied natural gas (LNG) imports to play a major role in bridging this demand-supply gap during the next several years. A global LNG market is quickly emerging, with several significant development projects underway. We present a practical real-option model for the valuation of downstream LNG storage. Our approach is based on a tractable stochastic dynamic-programming model to determine an optimal release-policy of downstream inventory released from a regasification terminal into the wholesale natural-gas market. This model uses a closed queuing network to represent upstream LNG production and shipping to the downstream regasification facility, and a reduced form model of the evolution of the natural-gas spot price.

Using our model, we show that the optimal LNG inventory-management policy at the downstream facility has a basestock-target structure that depends on the realization, at the inventory-review time, of the random state variable used to model the evolution of the natural-gas spot price. Contrary to what has been claimed by some practitioners, the structure of this policy is nontrivial; when it is optimal to sell it is not necessarily optimal to sell the entire available inventory. We also quantify the value of the real option to store LNG at this facility, and its relative benefit for different parties involved in an LNG value chain. While the storage option is equally valuable for LNG merchants and integrated producers, we find that its relative benefit is significantly higher for the former than the latter parties.

We also study the importance of modeling the shipping variability when valuing and managing a downstream LNG storage facility. The shipping model presented in Chapter 3 assumes that the unloaded amount in each decision period is independently and identically distributed (i.i.d.). We study the merit of the i.i.d. assumption by using simulation and developing an upper bound. We show that the model, that uses the i.i.d. assumption, provides a good estimation of the storage value, and yields a near optimal inventory control policy. We also test the performance of a model that uses constant throughput to determine the inventory release policy. This model performs worse than the model of Chapter 3 for storage valuation purposes, but can be used to suggest the optimal inventory control policy, especially when the ratio of flow rate to storage size is high, i.e., storage is scarce.

In our LNG storage valuation model, we assume the natural gas spot price and LNG import amount are independent. Current LNG imports in U.S. only account for around 2% of total gas consumption. Thus, LNG imports currently have small impact on U.S. natural

gas prices. In the future, U.S. LNG imports are expected to grow significantly (EIA 2006, [19]), and it will be interesting to incorporate their likely larger impact on U.S. natural gas prices when valuing and managing a U.S. bound LNG supply chain.

Bibliography

[1] S. Asmussen. *Applied Probability and Queues*. Springer, New York, NY, USA, 2003.

[2] R. Baldick, S. Kolos, and S. Tompaidis. Interruptible electricity contracts from an electricity retailer's point of view: Valuation and optimal interruption. *Operations Research*, 54:627–642, 2006.

[3] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22:248–260, 1975.

[4] F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179, 1976.

[5] P. P. Boyle and Y. K. Tse. An algorithm for computing values of options on the maximum or minimum of several assets. *Journal of Financial and Quantitative Analysis*, 25:215–227, 1990.

[6] M. Broadie and J. B. Detemple. Option pricing: Valuation models and applications. *Management Science*, 50:1145–1177, 2004.

[7] M. Broadie and P. Glasserman. Estimating security price derivatives using simulation. *Management Science*, 42:269–285, 1996.

[8] R. Caldentey, R. Epstein, and D. Sauré. Optimal exploitation of a nonrenewable resource. Working Paper, Stern School of Business, New York University, New York, NY, USA, 2007.

[9] R. Carmona and V. Durrleman. Pricing and hedging spread options. *Society for Industrial and Applied Mathematics Review*, 45:627–685, 2003.

[10] J. H. Cho, H. Kotzot, F. de la Vega, and C. Durr. Large lng carrier poses economic advantages, technical challenges. *LNG Observer*, 2, 2005.

[11] L. Clewlow and C. Strickland. *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, London, UK, 2000.

[12] J. W. Cohen and M. Rubinovitch. On level crossings and cycles in dam processes. *Mathematics of Operations Research*, 2:297–310, 1977.

[13] S. J. Deng, B. Johnson, and A. Sogomonian. Exotic electricity options and the valuation of electricity generation and transmission assets. *Decision Support Systems*, 30:383–392, 2001.

[14] D. Denny. Gtl: An attractive route for gas monetization? *Journal of Petroleum Technology Online*, 2005.

[15] D. Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ, USA, 2001.

[16] P. Enders, A. Scheller-Wolf, and N. Secomandi. Valuation of technology and extraction scaling options in natural gas production. Tepper Working Paper 2008-E04, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA, 2008.

[17] Energy Information Administration. U.S. Natural Gas Markets: Mid-Term Prospects for Natural Gas Supply. U.S. Department of Energy, Washington, DC, USA, 2001.

[18] Energy Information Administration. The Global Liquefied Natural Gas Market: Status & Outlook. U.S. Department of Energy, Washington, DC, USA, 2003.

[19] Energy Information Administration. Annual Energy Outlook 2006 with Projections to 2030. U.S. Department of Energy, Washington, DC, USA, 2006.

[20] Energy Information Administration. Natural Gas, in International Energy Outlook 2006. U.S. Department of Energy, Washington, DC, USA, 2006.

[21] A. Eydeland and K. Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons. Inc., New Jersey, USA, 2003.

[22] M. J. Faddy. Optimal control of finite dams: Continuous output procedure. *Advances in Applied Probability*, 6:689–710, 1975.

[23] A. R. Flower. Lng project feasibility. In G. B. Greenwald, editor, *Liquefied Natural Gas: Developing and Financing International Projects*. Kluwer Law International, London, UK, 1998.

[24] H. Geman. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, Chichester, UK, 2005.

[25] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, NY, USA, 2004.

[26] W. J. Hahn and J. S. Dyer. Discrete time modeling of mean-reverting stochastic processes for real option valuation. *European Journal of Operational Research*, 184:534–548, 2008.

[27] A. M. Hasofer. The almost full dam with poisson input. *Journal of the Royal Statistical Society, Series B*, 28:329–335, 1966.

[28] A. M. Hasofer. The almost full dam with poisson input: Further results. *Journal of the Royal Statistical Society, Series B*, 28:448–455, 1966.

[29] J. H. Holcomb. Storage downstream of regasification optimizes lng's value. *LNG Observer*, 2006.

[30] J. C. Hull. *Options, Futures, and Other Derivatives Securities*. Prentice Hall, Englewood Cliffs, NJ, USA, 2000.

[31] P. Jaillet, E. I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management Science*, 50:909–921, 2004.

[32] J. T. Jensen. The lng revolution. *The Energy Journal*, 24:1–45, 2003.

[33] H. Johnson. Options on the maximum or minimum of several assets. *Journal of Financial and Quantitative Analysis*, 22:227–283, 1987.

[34] R. Kamat and S. S. Oren. Exotic options for interruptible electricity supply contracts. *Operations Research*, 50:835–850, 2002.

[35] M. Kaplan, R. C. Wentworth, and R. J. Hischier. Simulation and optimization of lng shipping systems. In *ASME Transaction Petroleum Mechanical Engineering and Pressure Vessels and Piping Conference*.

[36] J. Keppo. Pricing of electricity swing options. *Journal of Derivatives*, 11:26–43, 2004.

[37] E. Kirk. Correlation in the energy markets. In *Managing Energy Price Risk*. Risk Publications and Enron, London, UK, 1995.

[38] E. Koenigsberg and R. C. Lam. Cyclic queue models of fleet operations. *Operations Research*, 24:516–529, 1976.

[39] D. G. Luenberger. *Investment Science*. Oxford University Press, New York, NY, USA, 1998.

[40] M. Manoliu. Storage options valuation using multilevel trees and calendar spreads. *International Journal of Theoretical and Applied Finance*, 7:425–464, 2004.

[41] W. Margrabe. The value of an option to exchange one asset for another. *Journal of Finance*, 33:178–186, 1978.

[42] V. Martínez-de-Albéniz and J. M. V. Simón. A capacitated commodity trading model with market power. Working Paper, IESE Business School, University of Navarra, Barcelona, Spain, 2007.

[43] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70:583–601, 2002.

[44] P. P. Miller. The chain of lng project contracts. In G. B. Greenwald, editor, *Liquefied Natural Gas: Developing and Financing International Projects*. Kluwer Law International, London, UK, 1998.

[45] P. A. P. Moran. *The Theory of Storage*. Wiley & Sons, Inc., New York, NY, USA, 1959.

[46] T. Osogami and M. Harchol-Balter. A closed-form solution for mapping general distributions to minimal ph distributions. *Performance Evaluation*, 63:524–552, 2006.

[47] S. R. Pliska. A diffusion process model for the optimal operations of a reservoir system. *Journal of Applied Probability*, 12:859–863, 1975.

[48] S. A. Ross. Hedging long run commitments: Exercises in incomplete market pricing. *Economic Notes*, 26:385–420, 1997.

[49] E. S. Schwartz. Review of *Investment under Uncertainty* by a. k. dixit and r. s. pindyck. *Journal of Finance*, 49:1924–1928, 1994.

[50] E. S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52:923–973, 1997.

[51] E. S. Schwartz and J. E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46:893–911, 2000.

[52] N. Secomandi. On the pricing of natural gas pipeline capacity. Tepper Working Paper 2004-E05, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.

[53] N. Secomandi. Optimal commodity trading with a capacitated storage asset. Tepper Working Paper 2005-E59, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.

[54] D. Seppi. Risk-neutral stochastic processes for commodity derivative pricing: An introduction and survey. In E. Ronn, editor, *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 2002.

[55] S. E. Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer, New York, NY, USA, 2004.

[56] J. E. Smith. Alternative approaches for solving real-options problems. *Decisions Analysis*, 2:89–102, 2005.

[57] J. E. Smith and K. F. McCardle. Valuing oil properties: Integrating options pricing and decision analysis approaches. *Operations Research*, 46:198–217, 1998.

[58] J. E. Smith and K. F. McCardle. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research*, 47:1–15, 1999.

[59] J. E. Smith and R. F. Nau. Valuing risky projects: Options pricing theory and decision analysis. *Management Science*, 41:795–816, 1995.

[60] R. M. Stulz. Options on the minimum or maximum of two risky assets. *Journal of Financial Economics*, 10:161–185, 1982.

[61] D. M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ, USA, 1998.

[62] C. L. Tseng and G. Barz. Short-term generation asset valuation: A real options approach. *Operations Research*, 50:297–310, 2002.

[63] C. L. Tseng and K. Y. Lin. A framework using two-factor price lattices for generation asset valuation. *Operations Research*, 55:234–251, 2007.

[64] M. X. Wang, N. Secomandi, S. Kekre, and A. Scheller-Wolf. Valuation of downstream liquefied-natural-gas storage. Tepper Working Paper 2006-E99, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA, 2007.