Tradeoff between Storage and Transport in Merchant Energy Trading on a Network

Selvaprabu Nadarajah

Liautaud Graduate School of Business, University of Illinois at Chicago, 601 South Morgan Street,
Chicago, IL, 60607, USA
selvan@uic.edu

Nicola Secomandi

Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA
15213-3890, USA
ns7@andrew.cmu.edu

Abstract

The operations of merchant energy trading in wholesale markets across different locations and current and future dates can be represented as a network where storage and transport trades compete for the capacity of storage and transport assets. We study the tradeoff between storage and transport trading for a network with a single storage asset and multiple transport assets, a realistic situation that we model as a Markov decision problem (MDP). Due to the intractability of computing an optimal policy of this MDP, we leverage our structural analysis of this model to modify a least squares Monte Carlo method to obtain a heuristic policy, also computing both lower and upper bounds on the market value of an optimal policy. On a realistic natural gas application, we document a substantial tradeoff between storage and transport trading. This tradeoff is difficult to manage, as sequential storage and transport trading is considerably suboptimal, especially when prioritizing transport over storage. In contrast, our joint policy is near optimal. A practice-based method based on sequentially reoptimizing a deterministic model is also near optimal, but, even after simplification, is computationally more intensive than our approach. Moreover, we highlight the operational differences between managing storage jointly with transport assets versus as a single asset. Beyond natural gas, our research has relevance for managing the merchant trading operations of other energy sources, natural resources, and other storable commodities.

1   Introduction

Energy plays an important economic role. For example, natural gas served more than one quarter of the 2012 energy consumption in the United States (EIA 2013). The availability and importance of natural gas is growing with the shale boom (Smith 2013). It is projected that natural gas consumption in North America will increase by 18% between 2008 and 2030 and be accompanied by a need for 130-210 billion US dollars worth of midstream natural gas infrastructure (INGAA 2009). Eighty percent of this projected infrastructure cost is for building new natural gas pipeline systems (INGAA 2009).

Being a commodity, energy is traded in wholesale markets (Eydeland and Wolyniec 2003, Geman 2005, and Secomandi and Seppi 2014). Specifically, merchants trade energy across different locations and current and future dates to take advantage of positive price differentials. The operations of

merchant energy trading can be modeled as a network where transport and storage trades compete for the finite capacity of storage and transport assets. The magnitude and management of the tradeoff between storage and transport trading in this setting is not well understood in the extant literature, which has so far studied the management of these activities in isolation. For example, work on natural gas storage includes Charnes et al. (1966), Lai et al. (2010), Secomandi (2010b), Boogert and De Jong (2011/12), Thompson (2012), Wu et al. (2012), Boogert and Mazières (2013), Nadarajah et al. (2014a), and Nadarajah et al. (2014b). Deng et al. (2001), Secomandi (2010a), and Secomandi and Wang (2012) consider energy transport. The business problem that we consider is thus novel.

We examine the tradeoff between storage and transport trading by considering a network that consists of a single storage asset and several transport assets, a realistic setting. We formulate a Markov decision problem (MDP) that models the trading of a merchant on such a network during a finite horizon. In every stage, the states of this MDP include the inventory of energy available in storage and the energy forward curves of a set of geographically interconnected wholesale energy markets – a forward curve is a vector of futures prices (Luenberger 1998, page 278). At each stage and state, the MDP multidimensional action is a vector of storage and transport trade amounts. Our model extends a growing literature on the real option (Dixit and Pindyck 1994) management of commodity conversion assets (Secomandi and Seppi 2014) by combining the management of storage and transport assets.

Computing an optimal policy of our MDP is intractable. Based on our analysis of this model, we obtain a heuristic policy and a lower bound on the market value of an optimal policy, modifying the version of the least squares Monte Carlo (LSM) approach (Longstaff and Schwartz 2001, Tsitsiklis and Van Roy 2001, Glasserman 2004, Chapter 8) studied by Nadarajah et al. (2014a). Specifically, we approximately solve a stochastic dynamic program (SDP) that is equivalent to our MDP. Here, we exploit our basestock target characterization of the storage component of an optimal policy of our model, a result that extends the known optimal storage policy structure in the absence of competing transport trades (Secomandi 2010b, Secomandi et al. 2014). We use our value function approximation to estimate a dual upper bound on the market value of an optimal policy (see Rogers 2002, Brown et al. 2010, and references therein).

We apply our method in the context of natural gas trading. In this setting, the network assets are

Electronic copy available at: http://ssrn.com/abstract=2483061

contracts that give merchant access to the capacity of interconnected pipelines and a storage facility. Specifically, the merchant owns the natural gas that pipeline companies transport or store on the merchant account. This contractual system describes the status quo of the natural gas industry in the Unites States. In particular, we focus on *firm* contracts that give merchants *guaranteed* access to natural gas storage and pipeline-transport capacity (Sturm 1997).

We consider a set of realistic natural gas instances developed in conjunction with an international energy trading company. We document a considerable tradeoff between storage and transport trading. This tradeoff is difficult to manage because sequential policies that favor one activity over the other are substantially suboptimal, especially when transport is given priority over storage. These insights appear novel. Our LSM-based heuristic policy is near optimal, a finding consistent with the results of Nadarajah et al. (2014a) in the absence of transport. Also near optimal is a version of the practice-based rolling intrinsic storage policy (Lai et al. 2010, Secomandi 2010b, Wu et al. 2012, Secomandi 2014) extended to deal with both storage and transport – this method is available in the FEA (2013) software. However, even after simplification, this policy is one order of magnitude slower to evaluate than our policy. The near optimality of this practice-based method is consistent with its performance observed for the no transport case (Lai et al. 2010, Secomandi 2014), but it is not obvious that such performance should persist in our network setting.

Operationally, managing storage in our natural gas network application differs from managing storage in the absence of competing transport assets. Specifically, compared to the latter case the storage flow rate and the average inventory are, respectively, considerably and slightly smaller in this network setting, and, via Little's law on average, purchased natural gas spends a longer time in storage before being sold back to the market, that is, the average flow time is longer. Although both this smaller flow rate and this longer average flow time are natural consequences of the substantial tradeoff between storage and transport trading, this smaller average inventory is less obvious and, again by Little's law, can be explained by the reduction in flow rate being stronger than the increase in average flow time. This analysis provides novel insights relative to the operational analysis of Secomandi et al. (2014) for the single storage asset. Moreover, the average storage per-unit margin is slightly smaller when storage and transport trades compete than when there are no competing transport trades, again due to the tradeoff between these trading activities.

Our research provides us with an improved understanding of the tradeoff between storage and

transport trading on a network, as well as of the differences between the joint merchant management of storage and transport in a network setting compared to managing storage in isolation. Moreover, our work offers merchants a more efficient method to near optimally manage storage in a network than is possible with our improved modifications of a method currently available in practice, as well as a way to assess the suboptimality of this and other heuristics.

Beyond natural gas, our research has relevance for managing the merchant trading operations of other energy sources, such as coal, electricity, and oil and petroleum products, natural resources, such as water and timber, and other storable commodities, such as agricultural products and metals (Markland 1975, Markland and Newett 1976, Smith and McCardle 1998, 1999, Deng et al. 2001, Kleindorfer and Wu 2003, Rømo et al. 2009, Löhndorf and Minner 2010, Boyabatli 2011, Boyabatli et al. 2011, Devalkar et al. 2011, Kim and Powell 2011, Lai et al. 2011, Arvesen et al. 2013, Löhndorf et al. 2013, Zhou et al. 2013, 2014, Salas and Powell 2014, Scott et al. 2014, Secomandi and Seppi 2014).

In §2 we provide a context for our model by discussing the trading of natural gas. In §3 we introduce our MDP and formalize the tradeoff between storage and transport trading. In §4 we reformulate our MDP as an SDP and analyze the value function and the storage component of an optimal policy of this SDP. In §5 we discuss our LSM-based policy and how to use it to estimate a lower bound on the combined market value of the storage and transport assets. In §6 we conduct our numerical analysis. We conclude in §7. Material in support of the discussion in §2 are in Online Appendix A. Proofs are in Online Appendix B. The estimation of a dual upper bound on the market value of an optimal policy of our MDP is discussed in Online Appendix C. Numerical results supplementing those in §6 are in Online Appendix D.

2 A Specific Context: Natural Gas Trading

Our model, presented in §3, is not specific to a given energy source. However, in this section we discuss the merchant trading of natural gas to provide a concrete context for our formulation.

Natural gas pipeline systems comprise of storage facilities, compressor stations, metering stations, and interconnect stations that link different pipelines (Pipeline Knowledge & Development 2010). Figure 1 illustrates the connections of the Bobcat storage facility, located in Louisiana, to five major pipeline systems owned and operated by the Texas Eastern Transmission Company
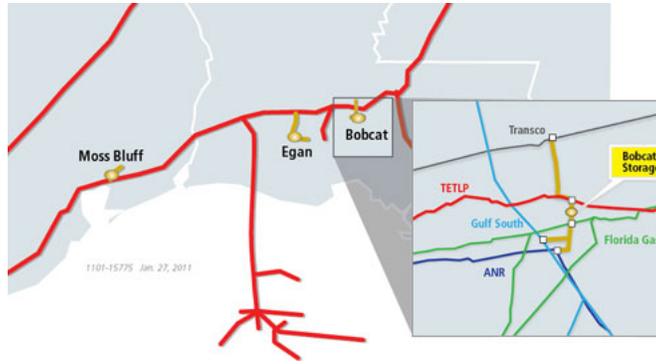
Figure 1: The Bobcat storage facility and its interconnecting pipelines (Source: Spectra Energy website).

(TETCO; also referred to as TETLP), the Transcontinental Gas Pipeline Company (TRANSCO), the Gulf South Pipeline Company, the Florida Gas Transmission Company, and the ANR Pipeline Company. Natural gas can be transferred across different pipelines through *interconnect stations*. Figure 1 shows that the Bobcat storage facility is an off-pipeline interconnect station. In contrast, the TETCO and Algonquin Gas Transmission (AGT) pipelines are directly connected at on-pipeline interconnect stations on the AGT pipeline (see Figure 9 in Online Appendix A).

Merchant trading of natural gas occurs on *commercial* networks, which are simplified representations of the physical pipeline systems that aggregate into *zones* pipeline segments, storage facilities, and compressor and metering stations. Figure 2 displays the zones of the TETCO pipeline. The zones of the TRANSCO and AGT pipelines are illustrated in Figures 10 and 11 in Online Appendix A. The AGT pipeline is treated as a single zone for merchant trading purposes due to its smaller size compared to both the TRANSCO and TETCO pipelines.

In North America, natural gas is traded on more than one hundred physical markets associated with pipeline zones. Financial derivative contracts on natural gas are traded on organized exchanges. Prominent examples are the New York Mercantile Exchange (NYMEX) natural gas futures contracts with delivery at Henry Hub, Louisiana, and basis swaps for about forty locations in North America – the price of a basis swap for a given maturity represents an offset with respect to the NYMEX natural gas futures price for the same maturity, and hence the futures price for such location and maturity is the sum of its corresponding basis swap price and the NYMEX futures price for this maturity. The NYMEX futures and basis swaps are associated with the zones of major North

American pipelines.

Figure 2: The TETCO pipeline system (Source: Rextag Strategies website).

The trading activity of merchants on natural gas commercial networks is based on acquiring contracts on the storage and transport capacity of pipelines (Sturm 1997). A storage contract specifies a collection of time periods during which storage can be used; the storage space accessible at a storage facility; injection and withdrawal capacities for each time period; and variable and fuel costs. A transport contract specifies a collection of time periods during which transport can be performed; a set of points where natural gas can be received into the pipeline (receipt points) or delivered from the pipeline (delivery points); capacity limits at each of these points; and variable and fuel costs to ship natural gas from receipt to delivery points. Commercially, the transport of natural gas is contemporaneous because natural gas is shipped by displacement using compressor stations that maintain pressure differentials between pipeline segments. We consider firm contracts, which give their owners guaranteed access to the reserved pipeline capacity and are associated with liquidated damages in case of pipeline nonperformance (NAESB 2002). We refer to these contracts as storage and transport *assets*. Merchants manage these assets as real options on natural gas prices that give them the ability to change the temporal or geographical availability of natural gas (Maragos 2002, Lai et al. 2010, Secomandi 2010a,b, Secomandi and Wang 2012).

6

3 Model and Tradeoff

We formulate our MDP in §3.1 and discuss the tradeoff in this model in §3.2.

3.1 Model

A merchant owns given energy transport and storage assets. We represent the location corresponding to the storage asset and the geographical markets connected by the transport assets as nodes on a network. These markets are included in the set $\mathcal{M}$. Although storage may be collocated with one of these markets, we represent it as a separate node to be able to distinguish between storage trades and transport trades, which we define below. Labeling storage by ST, the node set of the commercial network is $\mathcal{M} \cup \{ST\}$.

Figure 3 illustrates this representation using a realistic natural gas commercial network that includes (from left to right) three markets corresponding to TRANSCO zones 3, 4, and 6 (see Figure 10 in Online Appendix A), which we label Z3, Z4, and Z6; the Bobcat storage facility located at the interconnect station (IC) between TRANSCO and TETCO (see Figure 1); four markets on TETCO, corresponding to its zones 1 through 3 and East Louisiana (see Figure 2), which we label M1, M2, and M3, and ELA; and the AGT zone (see Figure 11 in Online Appendix A). An edge linking two nodes in this network indicates the possibility of transporting natural gas between these nodes (in both directions) using a transport asset (contract).
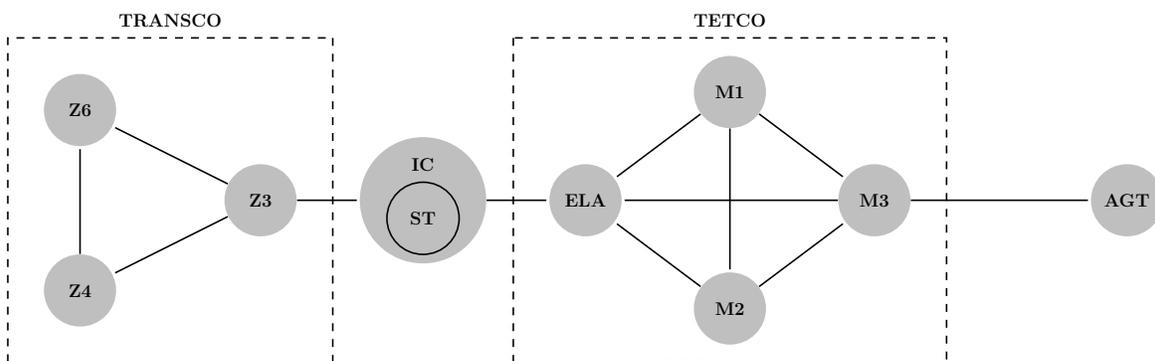
Figure 3: Natural gas commercial network including the Bobcat storage facility (ST), the interconnect station (IC), the subsets of the TRANSCO and TETCO zones (Z1-Z3, M1-M3, and ELA), and AGT.

The storage and transport assets allow the merchant to trade energy across different markets and dates. We define a merchant trade as an object that specifies the type of activity (storage or

transport), a date, and a unique path (sequence of nodes) in the commercial network. Our model is formulated based on the set of merchant trades, denoted by $\mathcal{J}$. The subsets $\mathcal{J}^I$ and $\mathcal{J}^W$ of $\mathcal{J}$ include the injection and withdrawal storage trades, respectively. We assume it is feasible to enumerate all the merchant trades supported by the merchant's transport and storage assets. If the cardinality of this set is too large, then we can easily reformulate our model without listing all the possible merchant trades by using instead the set of node-to-node flows as modeling objects.

We denote by p_j the path of trade j, and by $p_j(n)$ the n-th node in this path. The number of nodes in path p_j is denoted by $|p_j|$. In Figure 3, the paths Z3-Z6, M2-M3, Z3-IC-ELA, and M3-AGT belong to transport trades. The path of an injection trade ends at ST. The path of a withdrawal trade starts at ST. In Figure 3, paths Z3-IC-ST and AGT-M3-ELA-IC-ST belong to storage injection trades and paths ST-IC-Z3-Z6 and ST-IC-ELA to storage withdrawal trades.

Trades can be performed at each of N times. Denote the i-th trading time as T_i with i belonging to the set $\mathcal{I} := \{0, 1, \ldots, N-1\}$. We use the set $\mathcal{I}$ as the stage set of our MDP. Let $\bar{y}$ denote the maximum inventory allowed in the storage asset. The inventory in storage at stage i is $y_i \in \mathcal{Y} := [0, \bar{y}]$. We assume that a futures market is available at each market m. At time T_i the market m futures price with maturity at time $T_{i'} > T_i$ is $F_{i,i'}^m \in \mathbb{R}_+$ and the forward curve of this market is $\boldsymbol{F}_i^m := (F_{i,i}^m, F_{i,i+1}^m, \ldots, F_{i,N-1}^m)$. We denote the time T_i spot price at market m by $s_i^m \equiv F_{ii}^m$. We define the array of forward curves and the vector of spot prices across all markets at time T_i as $\boldsymbol{F}_i := (\boldsymbol{F}_i^m, \forall m \in \mathcal{M})$ and $\boldsymbol{s}_i := (s_i^m, \forall m \in \mathcal{M})$, respectively. We also define $\boldsymbol{F}_N := 0$. The stage i state is the pair $(y_i, \boldsymbol{F}_i)$.

The cash flow of a trade includes the cost of purchasing or revenue from selling energy at the prevailing spot price of the market where energy is transacted and two types of variables costs: Marginal costs and in-kind losses. The merchant incurs a marginal cost on each unit of energy transported, injected, or withdrawn. In-kind losses can model the use of energy to fuel the transport of energy in between nodes or the injection or withdrawal of energy into or out of storage, as well as inefficiencies when transporting energy or modifying the energy inventory. For example, compressor stations create pressure differentials between natural gas pipeline segments, enabling the transport of natural gas. Natural gas storage injections and withdrawals are also based on pressure differentials obtained by the use of pumps. Merchants pay the pipeline company in kind for the fuel used for compression.

The marginal cost for transporting energy between node m and node m' is denoted by $c^{m,m'}$. The storage injection and withdrawal marginal costs are denoted by c^I and c^W, respectively (I and W abbreviate injection and withdrawal, respectively). The in-kind loss $(1 - \phi^{m,m'})/\phi^{m,m'}$, where $\phi^{m,m'} \in (0,1]$, occurs when transporting 1 unit of energy from node m to node m': $1/\phi^{m,m'} \equiv 1 + (1 - \phi^{m,m'})/\phi^{m,m'}$ units of energy need to be received at node m in order to deliver 1 unit of energy at node m'. We assume that this energy is purchased at the market corresponding to node m. The in-kind losses incurred to inject and withdraw, respectively, 1 unit of energy into and out of storage are $(1 - \phi^I)/\phi^I$ and $(1 - \phi^W)/\phi^W$ where $\phi^I \in (0,1]$ and $\phi^W \in (0,1]$ have interpretations analogous to the transport in-kind loss. We assume that the energy used for injection or withdrawal is monetized at the spot price of the market closest to storage.

Denote by x_j the amount of energy transacted under trade $j \in \mathcal{J}$. We define the vector of trade amounts as $\boldsymbol{x} := (x_j, j \in \mathcal{J})$. The reward $r(\boldsymbol{x}, \boldsymbol{s})$ from executing the vector of trade amounts $\boldsymbol{x}$ given the vector of spot prices $\boldsymbol{s}$ is defined as

$$r(\boldsymbol{x}, \boldsymbol{s}) := \sum_{j \in \mathcal{J}} \sum_{n=1}^{|p_j|} \left[\alpha_n^j(\boldsymbol{s}) + \gamma_n^j \right] x_j, \tag{1}$$

where

$$\alpha_n^j(\boldsymbol{s}) := \begin{cases} -s^{p_j(2)} \dfrac{(1 - \phi^W)}{\phi^W} \mathbb{1}(j \in \mathcal{J}^W) - \dfrac{s^{p_j(1)}}{\phi^{p_j(1), p_j(2)}} \mathbb{1}(j \in \mathcal{J} \setminus \mathcal{J}^W), & \text{if } n = 1, \\[3mm] s^{p_j(l)} \dfrac{(1 - \phi^{p_j(n), p_j(n+1)})}{\phi^{p_j(n), p_j(n+1)}}, & \text{if } 1 < n < |p_j|, \\[3mm] -s^{p_j(|p_j|-1)} \dfrac{(1 - \phi^I)}{\phi^I} \mathbb{1}(j \in \mathcal{J}^I) + s^{p_j(|p_j|)} \mathbb{1}(j \in \mathcal{J} \setminus \mathcal{J}^I), & \text{if } n = |p_j|, \end{cases}$$

and

$$\gamma_n^j := \begin{cases} -c^W \mathbb{1}(j \in \mathcal{J}^W) - c^{p_j(1), p_j(2)} \mathbb{1}(j \in \mathcal{J} \setminus \mathcal{J}^W), & \text{if } n = 1, \\[2mm] -c^{p_j(n), p_j(n+1)}, & \text{if } 1 < n < |p_j|, \\[2mm] -c^I \mathbb{1}(j \in \mathcal{J}^I), & \text{if } n = |p_j|; \end{cases}$$

here, the term $\alpha_n^j(\boldsymbol{s})$ includes the per-unit cost incurred or revenue earned from buying or selling energy, respectively, and the corresponding monetized in-kind losses when executing trade j, and the term γ_n^j represents the marginal cost incurred when executing trade j.

We model capacity constraints by imposing limits on the maximum amount of energy that can be received or delivered at node m or that can be added or removed from storage during a single time period, that is, the time period elapsed in between two successive stages. The receipt and delivery capacities for node m are denoted as $C^{R,m}$ and $C^{D,m}$ (R is for receipt and D is for delivery), respectively. These capacity limits implicitly restrict the maximum amount of energy that can be transported between two nodes or between storage and a node during a single time period (Secomandi and Wang 2012). Natural gas storage and transport contracts (assets) specify node capacity limits. Alternatively, one could model capacity limits along each edge. The storage asset injection and withdrawal capacities are C^I and C^W, respectively.

Let $\mathcal{J}^R(m)$ and $\mathcal{J}^D(m)$ be the sets of trades with paths that include node m as a receipt point and a delivery point, respectively. We denote by $\vee$ a logical disjunction (or). A vector of trade amounts $\boldsymbol{x}$ is feasible at inventory level $y \in \mathcal{Y}$ if it satisfies the following conditions:

$$\sum_{j \in \mathcal{J}^R(m)} x_j \leq C^{R,m}, \forall m \in \mathcal{M}, \tag{2}$$

$$\sum_{j \in \mathcal{J}^D(m)} x_j \leq C^{D,m}, \forall m \in \mathcal{M}, \tag{3}$$

$$\left(\begin{array}{c} \sum_{j \in \mathcal{J}^I} x_j \leq \min\{C^I, \bar{y} - y\} \\ \sum_{j \in \mathcal{J}^W} x_j = 0 \end{array} \right) \vee \left(\begin{array}{c} \sum_{j \in \mathcal{J}^I} x_j = 0 \\ \sum_{j \in \mathcal{J}^W} x_j \leq \min\{C^W, y\} \end{array} \right), \tag{4}$$

$$x_j \geq 0, \forall j \in \mathcal{J}. \tag{5}$$

The node receipt and delivery capacities are imposed by constraints (2) and (3), respectively. The left and right hand sides of the disjunction (4) express the storage (flow) capacity constraints and the available space and inventory constraints: When the storage decision is to inject, (4) ensures that (i) the sum of the withdrawal trade amounts is zero and (ii) the sum of the injection trade amounts is less than both the storage injection capacity and the available space in storage; if the storage decision is to withdraw, (4) ensures that (i) the sum of the injection trade amounts is zero and (ii) the sum of the withdrawal trade amounts is less than both the storage withdrawal capacity and the available inventory in storage. Constraints (5) enforce nonnegativity of the trade amounts. The set of feasible trade amounts for feasible inventory level y is thus defined as $\mathcal{X}(y) := \{\boldsymbol{x} | (2)\text{-}(5)\}$.

10

Given the stage i spot price vector $\boldsymbol{s}_i$, executing a feasible collection of trade amounts $\boldsymbol{x}$ at inventory level y_i results in an immediate reward of $r(\boldsymbol{x}, \boldsymbol{s}_i)$ and an inventory transition from y_i to $y_i + \sum_{j \in \mathcal{J}^I} x_j - \sum_{j \in \mathcal{J}^W} x_j$.

The evolution of the stage i array of forward curves $\boldsymbol{F}_i$ into the stage $i{+}1$ array of forward curves $\boldsymbol{F}_{i+1}$ is governed by a known stochastic process, which is assumed to be unaffected by the merchant trades (that is, the merchant is a small player and, hence, a price taker). We model the continuous time *risk-neutral* dynamics of the array of forward curves using a multi-market version of the multifactor term structure model that is common both in the merchant energy trading literature and practice (Cortazar and Schwartz 1994, Clewlow and Strickland 2000, Blanco et al. 2002, Secomandi et al. 2014, and Secomandi and Seppi 2014, Chapter 4; the one and two factor models of Jaillet et al. 2004 and Schwartz and Smith 2000, respectively, as well as the multimarket specifications of the Jaillet et al. 2004 valuation model used by Secomandi 2010a and Secomandi and Wang 2012 are special cases of this model; Frestad 2008 and Suenaga et al. 2008 use higher dimensional models with maturity specific shocks). In this continuous time setting, we denote by $F^m(t, T_i)$ the market (node) m futures price at time $t \in [T_0, T_i]$ with maturity on date $T_i \geq t$. We let K be the number of stochastic factors driving the evolution of this price; $dW_k(t)$ the standard Brownian motion increment corresponding to factor k at time t (all these increments are uncorrelated, that is, $dW_k(t)dW_{k'}(t) = 0, \forall k, k' \in \{1, 2, \ldots, K\}, k \neq k'$); and $\sigma_{m,i,k}(t)$ the time t loading on the k-th factor for the market m futures with maturity at time T_i. The evolution of $F^m(t, T_i), \forall (m, i) \in \mathcal{M} \times \mathcal{I} \backslash \{0\}$ and $t \in (T_0, T_i]$, is governed by the following stochastic differential equations:

$$\frac{dF^m(t, T_i)}{F^m(t, T_i)} = \sum_{k=1}^{K} \sigma_{m,i,k}(t)dW_k(t), \ \forall (m, i) \in \mathcal{M} \times \mathcal{I} \backslash \{0\}, t \in (T_0, T_i]. \tag{6}$$

This model captures the seasonality in price levels via the initial (time T_0) array of forward curves, and the seasonalities in the price changes through the dependence of the loading factors on the trading time (t). The price changes are correlated because they are functions of common factors. Our analysis in §4 does not depend on this specific price model. In contrast, our algorithm developed in §5 to obtain a heuristic policy and the upper bound presented in Online Appendix C rely on this price model, of which we use a particular specification in our numerical analysis carried out in §6.

Model (6) is consistent with lack of arbitrage in futures markets. Under this model the realized

spot prices at two markets can differ by more than the marginal cost and in-kind loss incurred when transporting energy between these markets. An obvious modification of this statement involving the storage marginal costs and in-kind losses applies to the spot and futures prices for two different dates. These possibilities do not represent arbitrage opportunities; rather they are reduced form representations of what occurs in equilibrium models of energy prices on a network when a capacity constraint binds (see, e.g., Gabriel et al. 2005, §3.3, Secomandi 2010a, §5 and references therein).

Let $\mathbb{E}$ denote expectation under the corresponding risk-neutral probability measure for the forward curve evolution (Secomandi and Seppi 2014, Ch. 3), which is unique when the commodity market is complete (Björk 2004, page 122). Market completeness is a common assumption in the real option and the merchant trading literatures (see, for example, Smith and McCardle 1998, 1999, Lai et al. 2010, Devalkar et al. 2011). It holds in our setting when the number of traded futures contracts equals or exceeds the number of stochastic factors in model (6), the typical case in applications. A policy π is the collection of decision rules $\{A_0^\pi, A_1^\pi, \ldots, A_{N-1}^\pi\}$, where $A_i^\pi : (y_i, \boldsymbol{F}_i) \mapsto \mathcal{X}(y_i)$, $\forall (i, y_i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y} \times \mathbb{R}_+^{M \cdot (N-i)}$. We let Π be the set of all feasible policies. We denote by δ the risk-free discount factor from each time T_i back to time T_{i-1}, $\forall i \in \mathcal{I} \setminus \{0\}$. Let $(y_0, \boldsymbol{F}_0)$ be the time $T_0 := 0$ state. Maximizing the time T_0 market value of operating the storage and transport assets during the given time horizon entails solving the following MDP:

$$\max_{\pi \in \Pi} \sum_{i \in \mathcal{I}} \delta^i \mathbb{E}\left[r(A_i^\pi(y_i^\pi, \boldsymbol{F}_i), \boldsymbol{s}_i) | y_0, \boldsymbol{F}_0 \right], \tag{7}$$

where y_i^π is the random inventory level reached in stage i when using policy π.

3.2 Tradeoff

The main tradeoff in our MDP (7) is the competition between the storage and transport trades for the receipt and delivery capacity of the network nodes. Thus, intuitively, storage and transport trading are substitute activities. Proposition 3.1 formally states this property. Let Π^{TR} and Π^{ST} be the subsets of policies in Π that allow only transport and storage trades, respectively. We indicate by $V_0(y_0, \boldsymbol{F}_0)$ the optimal objective function value of model (7). We denote by $V_0^{TR}(y_0, \boldsymbol{F}_0)$ and $V_0^{ST}(y_0, \boldsymbol{F}_0)$ the optimal objective function values of the versions of this model with the restrictions $\pi \in \Pi^{TR}$ and $\pi \in \Pi^{ST}$, respectively.

Proposition 3.1. *It holds that $V_0(y_0, \boldsymbol{F}_0) \leq V_0^{TR}(y_0, \boldsymbol{F}_0) + V_0^{ST}(y_0, \boldsymbol{F}_0)$.*

The inequality in this proposition is consistent with the definition of substitutes in Topkis (1998, §2.6.1). When this inequality holds as a strict inequality, storage and transport trading are strict substitutes and jointly managing these activities is necessary to obtain an optimal policy: A pair of optimal policies to (7) subject to the restrictions $\pi \in \Pi^{TR}$ and $\pi \in \Pi^{ST}$, respectively, cannot form an optimal policy to (7) without these restrictions. When this inequality holds as an equality, there is no substitution between storage and transport trading, and hence they can be optimally managed independently of each other. We estimate numerically the degree of substitutability between these activities in §6.3.

4 Analysis

In §4.1 we reformulate our MDP (7) as an SDP to facilitate our analysis in §4.2 of the value function and the storage component of an optimal policy of this model. We use the results of this analysis in the development of our LSM approach in §5.

4.1 Reformulation of our MDP

Our starting point is the following equivalent reformulation of our MDP as an SDP, $\forall (i, y_i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y} \times \mathbb{R}_+^{M \cdot (N-i)}$:

$$V_i(y_i, \boldsymbol{F}_i) = \max_{\boldsymbol{x} \in \mathcal{X}(y_i)} r(\boldsymbol{x}, \boldsymbol{s}_i) + \delta \mathbb{E} \left[V_{i+1} \left(y_i + \sum_{j \in \mathcal{J}^I} x_j - \sum_{j \in \mathcal{J}^W} x_j, \boldsymbol{F}_{i+1} \right) \mid \boldsymbol{F}_i \right], \qquad (8)$$

with $V_i(y_i, \boldsymbol{F}_i)$ denoting the value function in stage i and state $(y_i, \boldsymbol{F}_i)$ and boundary conditions $V_N(y_N, \boldsymbol{F}_N) := 0, \forall y_N \in \mathcal{Y}$. Below we simplify this SDP by explicitly optimizing the inventory change assuming the storage and transport trading decisions are made optimally for every feasible inventory change.

We refer to the inventory change as the storage action, which, however, may result from executing more than one storage trade. Given the inventory levels y_i and y_{i+1} in stages i and $i+1$, respectively, define the storage action a as $y_i - y_{i+1}$: A negative (positive) storage action corresponds to an addition (removal) of energy from (to) the current inventory in the storage asset, where the sign of this action refers to the sign of the inventory addition cash flow, and a zero storage action

corresponds to leaving the inventory in storage unchanged (doing nothing). Let $\mathcal{X}'(a)$ denote the collections of all (storage and transport) trade amounts that satisfy the receipt and delivery capacity constraints and result in an inventory change equal to a. Given a storage action a, a vector of trade amounts $\boldsymbol{x}$ belongs to set $\mathcal{X}'(a)$ if it satisfies

$$\sum_{j \in \mathcal{J}^W} x_j = \begin{cases} a, & \text{if } a > 0, \\ 0, & \text{otherwise,} \end{cases} \tag{9}$$

$$\sum_{j \in \mathcal{J}^I} x_j = \begin{cases} -a, & \text{if } a < 0, \\ 0, & \text{otherwise,} \end{cases} \tag{10}$$

$$(2), (3), (5). \tag{11}$$

Constraints (9) and (10) ensure that the sums of the withdrawal and the injection trade amounts, respectively, are consistent with the given storage action a.

Given a spot price vector $\boldsymbol{s} \in \mathbb{R}^M_+$ and a storage action a, an optimal collection of storage and transport trade amounts in set $\mathcal{X}'(a)$ can be computed by solving the linear program

$$\bar{r}(a, \boldsymbol{s}) := \max_{\boldsymbol{x} \in \mathcal{X}'(a)} r(\boldsymbol{x}, \boldsymbol{s}). \tag{12}$$

We define a^I and a^W as the maximum injection and withdrawal amounts, respectively, ignoring the storage asset (flow) capacity constraints: $a^I := \max_{\boldsymbol{x}} \sum_{j \in \mathcal{J}^I} x_j$ s.t. (2), (3), (5), and $a^W := \max_{\boldsymbol{x}} \sum_{j \in \mathcal{J}^W} x_j$ s.t. (2), (3), (5). Given a storage action a, its feasibility, that is, $\mathcal{X}'(a) \neq \emptyset$, can be checked as follows: $\mathcal{X}'(a) \neq \emptyset$ if and only if $a \in [-a^I, a^W]$. Thus, the set of feasible stage $i+1$ inventory levels reachable from the stage i inventory level y_i is $\mathcal{Z}(y_i) := [\max\{0, y_i - C^W, y_i - a^W\}, \min\{\bar{y}, y_i + C^I, y_i + a^I\}]$.

Define the continuation function $W_i(y_{i+1}, \boldsymbol{F}_i)$, $\forall (i, y_{i+1}, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y} \times \mathbb{R}^{M \cdot (N-i)}_+$, as

$$W_i(y_{i+1}, \boldsymbol{F}_i) := \delta \mathbb{E}\left[V_{i+1}\left(y_{i+1}, \boldsymbol{F}_{i+1}\right) | \boldsymbol{F}_i\right]. \tag{13}$$

We can thus reformulate SDP (8) as

$$V_i(y_i, \boldsymbol{F}_i) = \max_{y_{i+1} \in \mathcal{Z}(y_i)} \bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i) + W_i(y_{i+1}, \boldsymbol{F}_i), \tag{14}$$

$\forall (i, y_i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y} \times \mathbb{R}_+^{M \cdot (N-i)}$, with boundary conditions $V_N(y_N, \boldsymbol{F}_N) := 0, \forall y_N \in \mathcal{Y}$. In contrast to SDP (8), the maximization in SDP (14) is over the feasible next-stage inventory level, and the function $\bar{r}(\cdot, \boldsymbol{s}_i)$ in the objective function of this maximization returns the value of the optimal storage and transport trade amounts conditional on a feasible storage action $y_i - y_{i+1}$.

4.2 Value and Continuation Functions and Structure of the Storage Component of an Optimal Policy

We characterize the value and continuation functions and under a mild assumption the structure of an optimal policy of SDP (14), that is, the *storage* component of an optimal policy. This analysis leads to a substantial simplification of SDP (14).

Lemma 4.1 states the concavity of the value and continuation functions of this SDP.

Lemma 4.1. *For each given $(i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathbb{R}_+^{M \cdot (N-i)}$, the functions $V_i(y_i, \boldsymbol{F}_i)$ and $W_i(y_{i+1}, \boldsymbol{F}_i)$ are concave in $y_i \in \mathcal{Y}$ and $y_{i+1} \in \mathcal{Y}$, respectively.*

We make a mild assumption that allows us to refine Lemma 4.1 and give the structure of the storage component of an optimal policy in Proposition 4.4.

Assumption 4.2. *The parameters $C^{R,m}, \forall m \in \mathcal{M}$, $C^{D,m}, \forall m \in \mathcal{M}$, y_0, $\bar{y}$, C^W, and C^I are rational numbers.*

Let $\bar{G}$ be the largest rational number such that the *transport-related* capacity values in sets $\{C^{R,m}, \forall m \in \mathcal{M}\}$ and $\{C^{D,m}, \forall m \in \mathcal{M}\}$ are integer multiples of $\bar{G}$ ($\bar{G}$ exists by Assumption 4.2). We interpret $\bar{G}$ as a lot size. Lemma 4.3 characterizes the function $\bar{r}(\cdot, \boldsymbol{s})$ in terms of this lot size.

Lemma 4.3. *Suppose Assumption 4.2 holds. Given the spot price vector $\boldsymbol{s} \in \mathbb{R}_+^M$, the function $\bar{r}(\cdot, \boldsymbol{s})$ is piecewise linear concave on the interval $[-a^I, a^W]$ with slope changes at integer multiples of $\bar{G}$.*

The structure of the storage component of an optimal policy given in Proposition 4.4 below relies on the *target* function $b_i(y_i, \boldsymbol{F}_i)$, defined as the smallest element of $\text{argmax}_{y_{i+1} \in \mathcal{Y}} \bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i) + W_i(y_{i+1}, \boldsymbol{F}_i)$. This function is qualified as a *target* function because it returns a stage $i+1$ inventory level that might not be reachable from the stage i inventory level y_i – the optimization in the definition of this function is over the set of all feasible inventory levels, $\mathcal{Y}$, which potentially *strictly*

includes the set $\mathcal{Z}(y_i)$ of feasible next-stage inventory levels for the given inventory level y_i. We define the functions $\underline{b}_i(\boldsymbol{F}_i)$ and $\bar{b}_i(\boldsymbol{F}_i)$ as those functions that give the smallest and largest feasible inventory levels, respectively, for which the do nothing storage action is optimal in stage i given the array of forward curves $\boldsymbol{F}_i$: $\underline{b}_i(\boldsymbol{F}_i) := \min_{y_i \in \mathcal{Y}} y_i$ s.t. $y_i = b_i(y_i, \boldsymbol{F}_i)$ and $\bar{b}_i(\boldsymbol{F}_i) := \max_{y_i \in \mathcal{Y}} y_i$ s.t. $y_i = b_i(y_i, \boldsymbol{F}_i)$. These functions exist, as stated in Proposition 4.4.

Under Assumption 4.2, Proposition 4.4, based on Lemmas 4.1 and 4.3, refines in Part (a) the characterizations of the value and continuation functions of SDP (14) given in Lemma 4.1 and characterizes in Part (b) the structure of the storage component of an optimal policy. This result depends on the lot size G, defined as the largest rational number such that $\bar{G}$, y_0, $\bar{y}$, C^W, and C^I are all integer multiples of G (G exists by Assumption 4.2). The lot size G is, potentially strictly, smaller than the lot size $\bar{G}$ because it also depends on the storage injection and withdrawal capacities and the initial (time T_0) and maximum inventory levels. We denote an optimal storage decision rule as $\bar{A}_i^*(y_i, \boldsymbol{F}_i)$.

Proposition 4.4. *Suppose Assumption 4.2 holds.*

(a) *Given $(i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathbb{R}_+^{M \cdot (N-i)}$, the value function $V_i(y_i, \boldsymbol{F}_i)$ and the continuation function $W_i(y_{i+1}, \boldsymbol{F}_i)$ are piecewise linear concave in $y_i \in \mathcal{Y}$ and $y_{i+1} \in \mathcal{Y}$, respectively, with slope changes at integer multiples of G.*

(b) *Given $(i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathbb{R}_+^{M \cdot (N-i)}$ and $q \in \{0, G, \ldots, (\bar{y}/G)-1\}$, the target function $b_i(\cdot, \boldsymbol{F}_i)$ equals a constant or varies linearly with slope 1 for all $y_i \in [qG, (q+1)G]$: $b_i(y_i, \boldsymbol{F}_i) = b_i(qG, \boldsymbol{F}_i) + (y_i - qG)\theta_q, \forall y_i \in [qG, (q+1)G]$, with $\theta_q \in \{0, 1\}$. The functions $\underline{b}_i(\cdot)$ and $\bar{b}_i(\cdot)$ exist and partition the feasible inventory set $\mathcal{Y}$ into the three regions $[0, \underline{b}_i(\boldsymbol{F}_i))$, $[\underline{b}_i(\boldsymbol{F}_i), \bar{b}_i(\boldsymbol{F}_i)]$, and $(\bar{b}_i(\boldsymbol{F}_i), \bar{y}]$, such that (i) an optimal storage action is to inject when $y_i \in [0, \underline{b}_i(\boldsymbol{F}_i))$, do nothing when $y_i \in [\underline{b}_i(\boldsymbol{F}_i), \bar{b}_i(\boldsymbol{F}_i)]$, and withdraw when $y_i \in (\bar{b}_i(\boldsymbol{F}_i), \bar{y}]$, and (ii) the function $b_i(\cdot, \boldsymbol{F}_i)$ evaluated at y_i returns a value that lies in the region of this partition that y_i belongs to. Specifically, the storage component of an optimal policy satisfies, $\forall (i, y_i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y} \times \mathbb{R}_+^{M \cdot (N-i)}$,*

$$\bar{A}_i^*(y_i, \boldsymbol{F}_i) = \begin{cases} \max\{-C^I, y_i - b_i(y_i, \boldsymbol{F}_i)\}, & \text{if } y_i \in [0, \underline{b}_i(\boldsymbol{F}_i)), \\ 0, & \text{if } y_i \in [\underline{b}_i(\boldsymbol{F}_i), \bar{b}_i(\boldsymbol{F}_i)], \\ \min\{C^W, y_i - b_i(y_i, \boldsymbol{F}_i)\}, & \text{if } y_i \in (\bar{b}_i(\boldsymbol{F}_i), \bar{y}]. \end{cases} \quad (15)$$
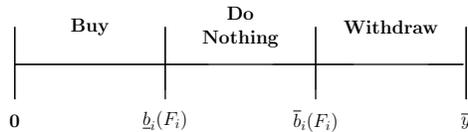
Figure 4: Partition of the feasible inventory set based on the type of storage action

Part (a) of Proposition 4.4 is related to results established by Bannister and Kaye (1991) and Nascimento and Powell (2013). Part (b) of Proposition 4.4 extends to the multiple market setting the known *double basestock* target structure of an optimal storage policy in the single market case (Secomandi 2010b, Secomandi et al. 2014).

Before interpreting our structure of the storage component of an optimal policy, we briefly describe the structure of an optimal storage policy that is known in the single market case. This structure includes two stage-and-forward-curve dependent basestock target functions. Given a stage and a forward curve, it is optimal to inject up to the smaller basestock target function value for inventory levels below this target value, withdraw down to the larger basestock target function value for inventory levels above this target value, and do nothing for inventory levels between these target values. Thus, for each given stage and forward curve, the feasible inventory set is partitioned into three contiguous inject, do nothing, and withdraw regions, and the optimal storage action within the inject and withdraw regions *strictly* increases in inventory. Moreover, under an assumption analogous to Assumption 4.2, but which excludes the initial inventory level and transport-related capacities, the latter of which are irrelevant in the single market case, these target values are integer multiples of a given lot size.

Analogous to the single market case, for each given stage and array of forward curves, in our case the feasible inventory set is partitioned into three contiguous inject, do nothing, and withdraw regions (see Figure 4). However, this partition is defined by the functions $\underline{b}_i(\boldsymbol{F}_i)$ and $\bar{b}_i(\boldsymbol{F}_i)$, which may not be target functions. In other words, in general these are only threshold functions that define this partition. Moreover, an optimal storage action is determined by the target function $b_i(y_i, \boldsymbol{F}_i)$, which can be interpreted as a piecewise linear *basestock* target function that also depends on the inventory level y_i, and can thus take *infinitely* many values for each given array of forward curves $\boldsymbol{F}_i$. Nevertheless, because the slope of the linear segment of the basestock target function is 1, there are only finitely many, and potentially more than two, *reachable* target values starting

from feasible inventory levels that are integer multiples of G. As in the single market case, however, these basestock target functions do not bring the resulting next-stage inventory level outside of the region a given current inventory level belongs to in the stated inject, do nothing, and withdraw partition, and the resulting next-stage inventory level is a weakly increasing function of the current inventory level. Moreover, different from the single market case, the optimal storage injection and withdrawal actions in the multiple market setting can be *weakly*, rather than strictly, increasing in the inventory level, that is, it is possible to have subregions where these actions are constant. Figure 5 conceptually illustrates this structure for the injection case.
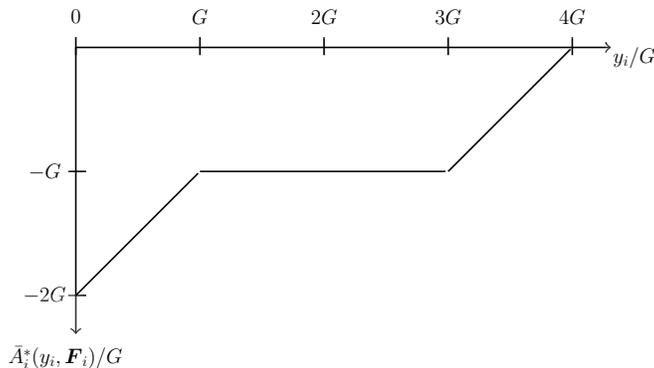
Figure 5: Conceptual illustration of the piecewise linearity of an optimal policy structure in the injection region with $\underline{b}_i(\boldsymbol{F}_i) = 4G$.

Our structure of the storage component of an optimal policy has useful *computational* implications. By Part (b) of Proposition 4.4, it is optimal to visit only inventory levels in the finite set $\mathcal{Y}' :=$ $\{0, G, \ldots, \bar{y}\}$, with $\bar{y}/G + 1$ values, and for each inventory level y_i in this set the next-stage inventory level is optimally chosen from the finite set $\mathcal{Z}'(y_i) := \{y_i - \min\{C^W, a^W\}, y_i - \min\{C^W, a^W\} + G, \ldots, y_i + \min\{C^I, a^I\}\}$. Thus, in principle, the structure of the optimal storage decision rules (15) could be used, as explained below, to solve the following SDP, $\forall (i, y_i, \boldsymbol{F}_i) \in \mathcal{I} \times \mathcal{Y}' \times \mathbb{R}_+^{M \cdot (N-i)}$:

$$V_i(y_i, \boldsymbol{F}_i) = \max_{y_{i+1} \in \mathcal{Z}'(y_i)} \bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i) + W_i(y_{i+1}, \boldsymbol{F}_i), \tag{16}$$

with boundary conditions $V_N(y_N, \boldsymbol{F}_N) := 0, \forall y_N \in \mathcal{Y}'$. SDP (16) critically differs from SDP (14) because it optimizes over the finite set $\mathcal{Z}'(y_i)$ rather than the interval $\mathcal{Z}(y_i)$ and has a finite number of inventory levels in its states.

18

Given the pair $(i, \boldsymbol{F}_i)$ and *assuming* knowledge of the function $W_i(\cdot, \boldsymbol{F}_i)$, the optimization on the right hand side of (16) can be performed efficiently *for all* the feasible inventory levels in set $\mathcal{Y}'$. Specifically, combining the structure of the functions $\bar{r}(\cdot, \boldsymbol{s}_i)$ and $W_i(\cdot, \boldsymbol{F}_i)$ and the decision rules (15) allows us to efficiently compute the set of optimal next-stage inventory levels for all inventory levels $y_i \in \mathcal{Y}'$ by scanning *once* the next-stage inventory levels in the set $\mathcal{Y}'$. This efficient search is possible because the optimal next-stage inventory level weakly increases in the current stage inventory level. This property allows us to restrict search for the optimal next-stage inventory level for a given current inventory level y_i in set $\mathcal{Y}' \setminus \{0\}$ to the subset of $\mathcal{Y}'$ delimited below by the optimal next-stage inventory level computed for $y_i - G$ and above by the maximum inventory level $\bar{y}$.

Although the function $W_i(\cdot, \boldsymbol{F}_i)$ is unknown, the same scheme just described remains applicable when this function is replaced by another piecewise linear concave function of the current inventory level with slope changes at integer multiples of the lot size G. We use this approach in §5 to develop our approximate solution method, which relies on approximating the unknown continuation function $W_i(\cdot, \boldsymbol{F}_i)$ with a function that satisfies this property. Moreover, we take advantage of this property when using our approximate continuation function to compute a heuristic policy and estimate a lower bound on the combined market value of the storage and transport assets, as detailed in §5. The approach presented in Online Appendix C to estimate a dual upper bound on this market value also critically relies on our value function approximations, which we use to obtain our approximate continuation functions, changing slope at integer multiples of the lot size G.

The lot size used to define the sets $\mathcal{Y}'$ and $\mathcal{Z}'(y_i)$ depends on the initial (time T_0) inventory level. This dependence implies that different initial inventory levels potentially require different discretizations of these sets. In contrast, the lot size used to optimally discretize the feasible inventory and storage action sets in the single market case is independent of the initial inventory level (Secomandi 2010b, Secomandi et al. 2014).

5 LSM-based Heuristic Policy and Lower Bound

In theory, an optimal action at a given stage and state can be computed by solving the optimization problem defined by the right hand side of SDP (16) for the optimal storage action – we suppose that Assumption 4.2 holds in the rest of this paper – which also yields the optimal transport decisions

from the associated optimal solution of the linear program (12). However, this approach is not practical because computing the continuation function in the storage optimization is intractable due to (i) the high dimensionality of the state space and (ii) the inability to compute the expectation in the definition of this function. To overcome both these issues, we follow a common approximate dynamic programming (ADP) approach to compute heuristic, but hopefully near optimal, decisions, by replacing the unknown continuation function on the right hand side of SDP (16) by a *tractable* continuation function approximation (Powell 2011, Bertsekas 2012). A lower bound on the market value of the storage and transport assets can be estimated by Monte Carlo simulation of the policy defined by these heuristic decisions. We explain the details of the lower bound estimation after we describe the LSM approach (Longstaff and Schwartz 2001, Tsitsiklis and Van Roy 2001) that we use to compute a continuation function approximation.

We extend the regress-later version of the LSM approach (see Nadarajah et al. 2014a and references therein) to compute a continuation function approximation by first computing a *value function* approximation, which is also remarkably useful for numerically estimating a dual upper bound on the value of the storage and transport assets, as discussed in Online Appendix C. We consider value function approximations that are linear combinations of a given set of basis functions. For each stage i and inventory level $y_i \in \mathcal{Y}'$ we specify B_i basis functions. The b-th basis function is $\phi_{i,y_i,b} : \boldsymbol{F}_i \mapsto \mathbb{R}$ and its linear combination weight is $\beta_{i,y_i,b}$. Defining the row and column vectors $\Phi_{i,y_i} := (\phi_{i,y_i,1}, \phi_{i,y_i,2}, \ldots, \phi_{i,y_i,B_i})$ and $\beta_{i,y_i} := (\beta_{i,y_i,1}, \beta_{i,y_i,2}, \ldots, \beta_{i,y_i,B_i})^\top$, the value function approximation is

$$\hat{V}_i(y_i, \boldsymbol{F}_i; \beta_{i,y_i}) := (\Phi_{i,y_i} \beta_{i,y_i})(\boldsymbol{F}_i) \equiv \sum_{b=1}^{B_i} \phi_{i,y_i,b}(\boldsymbol{F}_i) \beta_{i,y_i,b}. \tag{17}$$

We obtain a continuation function approximation by replacing V_i by $\hat{V}_i$ on the right hand side of (13):

$$\hat{W}_i(y_{i+1}, \boldsymbol{F}_i; \beta_{i+1,y_{i+1}}) = \delta \sum_{b=1}^{B_{i+1}} \mathbb{E}\left[\phi_{i+1,y_{i+1},b}(\boldsymbol{F}_{i+1}) | \boldsymbol{F}_i\right] \beta_{i+1,y_{i+1},b}. \tag{18}$$

Following Nadarajah et al. (2014a) we choose basis functions such that the expectations in (18) can be computed exactly when using the price model (6), as discussed in §6.2. We label modified LSM (MLSM) our modification of the LSM variant investigated by Nadarajah et al. (2014a). Our

modification consists of a concavification step, discussed below, which yields a continuation function approximation $\hat{W}_i^{CN}(\cdot, \boldsymbol{F}_i; \beta_{i+1,\cdot})$, where the superscript abbreviates concave, which is piecewise linear concave with break points in set $\mathcal{Y}'$ (see Powell 2011, Ch. 13 and Nascimento and Powell 2013 for the use of a similar approach in related ADP contexts).

Algorithm 1: MLSM

Inputs: Number of sample paths H and basis functions that allow the exact computation of the expectations in (18).

Outputs: Weights $\hat{\beta}_{i,y_i}$, $\forall (i, y_i) \in \mathcal{I} \times \mathcal{Y}'$.

Initialization: Generate H *regression* sample paths of the arrays of forward curves $\{\boldsymbol{F}_i^h, i \in \mathcal{I} \setminus \{0\}, h = 1, \ldots, H\}$ starting from $\boldsymbol{F}_0$; $\hat{\beta}_{N,y_N} := 0, \forall y_N \in \mathcal{Y}'$.

For each $i = N - 1$ to 1 **do**:

    1. **For** each $h \in \{1, 2, \ldots, H\}$ **do**:

        (i) **For** each $y_{i+1} \in \mathcal{Y}'$ **do**:

$$\hat{W}_i(y_{i+1}, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,y_{i+1}}) := \delta \sum_{b=1}^{B_{i+1}} \mathbb{E}\left[\phi_{i+1,y_{i+1},b}(\boldsymbol{F}_{i+1}) | \boldsymbol{F}_i^h\right] \hat{\beta}_{i+1,y_{i+1},b}.$$

        (ii) Concavify $\hat{W}_i(\cdot, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,\cdot})$ to obtain $\hat{W}_i^{CN}(\cdot, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,\cdot})$.

        (iii) **For** each $y_i \in \mathcal{Y}'$ **do**:

$$v_i(y_i, \boldsymbol{F}_i^h) := \max_{y_{i+1} \in \mathcal{Z}'(y_i)} \hat{r}(y_i - y_{i+1}, \boldsymbol{s}_i^h) + \hat{W}_i^{CN}(y_{i+1}, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,y_{i+1}}). \tag{19}$$

    2. **For** each $y_i \in \mathcal{Y}'$ **do**: Perform a 2-norm regression on the set of approximate value function estimates $\{v_i(y_i, \boldsymbol{F}_i^h), \forall h \in \{1, 2, \ldots, H\}\}$ to determine the weights $\hat{\beta}_{i,y_i}$.

Algorithm 1 summarizes MLSM. The inputs to MLSM are the number of sample paths and basis functions that allow the exact computation of the expectation in (18) (see §6.2 for an example). MLSM outputs the weights $\hat{\beta}_{i,y_i}$ that define a value function approximation via (17) and then yield a continuation function approximation via (18). MLSM starts by generating H *regression* sample paths of the array of forward curves from stage 1 through N starting from $\boldsymbol{F}_0$, which we include in set $\{\boldsymbol{F}_i^h, i \in \mathcal{I} \setminus \{0\}, h = 1, 2, \ldots, H\}$, and initializing the stage N weight vector $\hat{\beta}_{N,y_N}$ to zero. At each stage $i \in \mathcal{I} \setminus \{0\}$, starting from stage $N - 1$ and moving backward to stage 1: In Step 1(i), for each $h \in \{1, 2, \ldots, H\}$, MLSM computes the stage i continuation function approximation using the stage $i + 1$ basis function weights; in Step 1(ii) MLSM concavifies this continuation
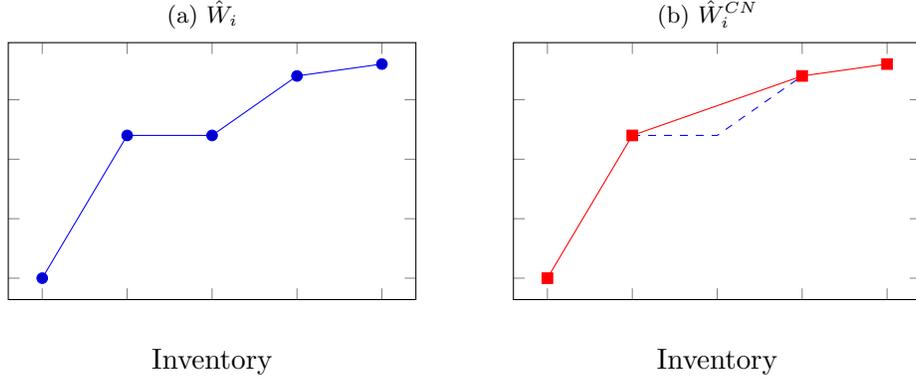
Figure 6: Illustration of the MLSM concavification step.

function approximation by a simple scanning and modification procedure illustrated in Figure 6; in Step 1(iii) MLSM computes the approximate value function estimates for each inventory level $y_i \in \mathcal{Y}'$ by solving a set of optimization problems (19) that are greedy with respect to the sum of the immediate reward function and the continuation function approximation – these optimizations can be performed efficiently using the scheme discussed at the end of §4.2, because the functions $\hat{W}_i^{CN}(y_{i+1}, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,y_{i+1}})$ and $\hat{r}(y_i - y_{i+1}, \boldsymbol{s}_i^h)$ are both piecewise linear concave in y_{i+1} with break points $\mathcal{Z}'(y_i)$ given $(i, y_i, \boldsymbol{F}_i^h)$; in Step 2, MLSM executes a 2-norm regression on these value function estimates to determine the weight vectors $\hat{\beta}_{i,y_i}$ at each inventory level y_i in set $\mathcal{Y}'$.

We now explain how the continuation function approximation estimated by MLSM is used to estimate a lower bound. At stage i and state $(y_i, \boldsymbol{F}_i)$, we replace the continuation function W_i on the right hand side of SDP (16) with $\hat{W}_i^{CN}(y_{i+1}, \boldsymbol{F}_i; \hat{\beta}_{i+1,y_{i+1}})$ to obtain the following optimization problem, which is greedy with respect to the sum of $\hat{r}$ and $\hat{W}_i^{CN}$:

$$\max_{y_{i+1} \in \mathcal{Z}'(y_i)} \hat{r}(y_i - y_{i+1}, \boldsymbol{s}_i) + \hat{W}_i^{CN}(y_{i+1}, \boldsymbol{F}_i; \hat{\beta}_{i+1,y_{i+1}}). \tag{20}$$

Given the solution y_{i+1}^g to (20), where the superscript abbreviates greedy and we break ties in favor of the smallest maximizer, the corresponding greedy storage action is $a_i^g := y_i - y_{i+1}^g$. The collection of greedy storage actions at all stages and states defines the greedy storage policy associated with $\hat{W}_i^{CN}$. Because $\hat{W}_i^{CN}(\cdot, \boldsymbol{F}_i^h; \hat{\beta}_{i+1,\cdot})$ is piecewise linear concave with slope changes at integer multiples of the lot size G, each y_{i+1}^g can be found efficiently, and the corresponding greedy storage policy has the same structure of the storage component of an optimal policy stated in Part (b) of Proposition

4.4. A lower bound on the value of an optimal policy can be estimated by applying the greedy storage policy and its associated transport policy, which we obtain via (12) with a_i^g in lieu of a, in Monte Carlo simulation. Specifically, we apply these policies along L Monte Carlo *simulation* sample paths of the array of forward curves, which we include in set $\{\boldsymbol{F}_i^l, i \in \mathcal{I} \setminus \{0\}, l = 1, 2, \ldots, L\}$, starting from the time 0 inventory level y_0 and the array of forward curves $\boldsymbol{F}_0$, and estimate a lower bound by averaging the resulting total discounted cash flows.

6 Natural Gas Application

In this section we discuss our natural gas application. In §6.1 we describe our instances. In §6.2 we estimate bounds on the combined market value of the storage and transport assets based on our MLSM approach. In §6.3 we analyze the tradeoff between storage and transport trading. In §6.4 we investigate the performance of the extended rolling intrinsic policy, and simplified versions thereof. In §6.5 we contrast the management of storage as a single asset versus jointly with transport assets.

6.1 Instances

We developed our instances in conjunction with a major multinational energy company that also operates in the United States. These instances are based on the commercial network displayed in Figure 3 and discussed in §3. We do not explicitly model the interconnect station IC, in addition to ST, because the marginal costs and in-kind losses for transporting natural gas between IC and ST are zero. We use a time horizon equal to one year subdivided into monthly periods (that is, $N = 12$).

The storage asset parameters are normalized maximum inventory ($\bar{y}$) equal to 1 MMBtu; normalized monthly injection capacity (C^I) and withdrawal capacity (C^W) equal to 0.45 MMBtu/month and 0.75 MMBtu/month, respectively; injection and withdrawal fuel losses (ϕ^I and ϕ^W) equal to 1 and 0.985, respectively; and injection and withdrawal commodity charges (c^I and c^W) equal to $0.02 /MMBtu and $0.01 /MMBtu, respectively (in the natural gas industry, an in-kind loss is known as a fuel loss and a marginal cost as a commodity charge).

The parameters of the transport assets are commodity charges and fuel losses ($c^{m,m'}$ and $\phi^{m,m'}$, respectively) as given in Tables 1-3; receipt and delivery capacities at both markets (nodes) Z3 and ELA equal to 0.45 MMBtu/month ($= C^W$) and 0.75 MMBtu/month ($= C^I$), respectively;

Table 1: Transport fuel losses ($\phi^{m,m'}$) for the months April to November.

	ST	Z3	Z4	Z6	ELA	M1	M2	M3	AGT
ST	-	1	-	-	1	-	-	-	-
Z3		-	0.9823	0.9638	-	-	-	-	-
Z4			-	0.9672	-	-	-	-	-
Z6				-	-	-	-	-	-
ELA					-	0.9557	0.9406	0.9305	-
M1						-	0.9632	0.9531	-
M2							-	0.9602	-
M3								-	0.9907

Table 2: Transport fuel losses ($\phi^{m,m'}$) for the months December to March.

	ST	Z3	Z4	Z6	ELA	M1	M2	M3	AGT
ST	-	1	-	-	1	-	-	-	-
Z3		-	0.9823	0.9638	-	-	-	-	-
Z4			-	0.9672	-	-	-	-	-
Z6				-	-	-	-	-	-
ELA					-	0.9523	0.9316	0.9179	-
M1						-	0.956	0.9423	-
M2							-	0.952	-
M3								-	0.99

Table 3: Commodity charges ($c^{m,m'}$, \$/MMBtu).

	ST	Z3	Z4	Z6	ELA	M1	M2	M3	AGT
ST	-	0.05	-	-	0.0103	-	-	-	-
Z3		-	0.02253	0.04454	-	-	-	-	-
Z4			-	0.04027	-	-	-	-	-
Z6				-	-	-	-	-	-
ELA					-	0.0353	0.0762	0.1044	-
M1						-	0.0659	0.0941	-
M2							-	0.0743	-
M3								-	0.013

receipt and delivery capacities at all the TRANSCO markets other than market Z3 equal to 0.15 MMBtu/month ($= C^W/3$) and 0.25 MMBtu/month ($= C^I/3$), respectively; and receipt and delivery capacities at the AGT market and all the TETCO markets other than market ELA equal to 0.09 MMBtu/month ($= C^W/5$) and 0.15 MMBtu/month ($= C^I/5$), respectively. Thus, the lot size G is equal to 0.01.

We calibrated a specification of the futures price model (6) using data made available to us by the energy trading company that we engaged. In this specification each function $\sigma_{m,i',k}(t)$ associated with maturity time $T_{i'}$ is right continuous and piecewise constant within a trading month, that is,

during each interval in the set $\{[T_i, T_{i+1}), \forall i \in \mathcal{I} \setminus \{N-1\}\}$ (Blanco et al. 2002 and Secomandi et al. 2014 use analogous specifications). We denote by $\sigma_{m,i',k,i}$ the value taken by each such function in the interval $[T_i, T_{i+1})$. Pick $t \in [T_i, T_{i+1})$ and $t' \in (T_i, T_{i+1}]$ with $t' > t$. Let $\mathcal{I}_i := \{i+1, i+2, \ldots, N-1\}$. We can thus express (6) in a form suitable for Monte Carlo simulation using K independent standard normal random variables, Z_k's, $\forall (m, i, i') \in \mathcal{M} \times \mathcal{I} \times \mathcal{I}_i$ as

$$F^m(t', T_{i'}) = F^m(t, T_{i'}) \exp\left[-\frac{1}{2}(t' - t)\sum_{k=1}^{K}\sigma^2_{m,i',k,i} + \sqrt{t' - t}\sum_{k=1}^{K}\sigma_{m,i',k,i}Z_k\right]. \qquad (21)$$

Our data set includes 1 year and 3 months of natural gas closing futures prices for Henry Hub, Louisiana, and basis swaps from June 2011 to August 2012 for each of the 8 markets in Figure 3, from which we created a futures price data set for these 8 markets. We first estimated monthly sample variance-covariance matrices of the daily log futures price returns across maturities and markets. We then performed a principal component analysis of these matrices and estimated the loading coefficients $\sigma_{m,i',k,i}$'s accordingly (see Blanco et al. 2002 and Secomandi et al. 2014 for details). We chose the number of factors K equal to 6 because this is the smallest value that explains more than 99% of the total observed variance in each of our monthly data sets.

We created 12 instances by choosing 12 valuation dates corresponding to the first trading date of each month from June 2011 to May 2012. We set the discount factor for each instance based on the following one year United States Treasury rates corresponding to our valuation dates: 0.18%, 0.20%, 0.22%, 0.10%, 0.12%, 0.13%, 0.12%, 0.12%, 0.13%, 0.18%, 0.18%, and 0.19%. The details of the estimated loading coefficients and the initial forward curves are available upon request.

6.2 MLSM-based Estimated Lower and Upper Bounds

For a given stage and inventory pair (i, y_i), we implemented MLSM using the following polynomials of futures prices as basis functions: 1, $\{F^m_{i,i'}, \forall i' \in \mathcal{I}_i\}$, $\{(F^m_{i,i'})^2, \forall i' \in \mathcal{I}_i\}$, $\{F^m_{i,i'}F^{m'}_{i,i'}, \forall i' \in \mathcal{I}_i; m, m' \in \mathcal{M}, m \neq m'\}$, and $\{F^m_{i,i'}F^m_{i,i'+1}, i' \in \mathcal{I}_i \setminus \{N-1\}, m \in \mathcal{M}\}$. This choice of basis functions is common in the LSM literature (Longstaff and Schwartz 2001) and Nadarajah et al. (2014a) use it for valuing storage in a single market. Under price model (21), defining $\Delta T_i := T_{i+1} - T_i$, it is easy to verify that for $i'' \geq i' > i$ the expectation in (18) for each of these basis functions is

$$\mathbb{E}[F^m_{i',i''} | F^m_{i,i''}] = F^m_{i,i''},$$

$$\mathbb{E}[(F_{i',i''}^m)^2 | F_{i,i''}^m] = (F_{i,i''}^m)^2 \exp\left\{\sum_{l=i}^{i'-1} \Delta T_l \sum_{k\in\mathcal{K}} \sigma_{m,i'',k,l}^2\right\},$$

$$\mathbb{E}[F_{i',i''}^m F_{i',i''}^{m'} | F_{i,i''}^m, F_{i,i''}^{m'}] = F_{i,i''}^m F_{i,i''}^{m'} \exp\left\{\sum_{l=i}^{i'-1} \Delta T_l \sum_{k\in\mathcal{K}} \sigma_{m,i'',k,l}\sigma_{m',i'',k,l}\right\},$$

$$\mathbb{E}[F_{i',i''}^m F_{i',i''+1}^m | F_{i,i''}^m, F_{i,i''+1}^m] = F_{i,i''}^m F_{i,i''+1}^m \exp\left\{\sum_{l=i}^{i'-1} \Delta T_l \sum_{k\in\mathcal{K}} \sigma_{m,i'',k,l}\sigma_{m,i''+1,k,l}\right\}.$$
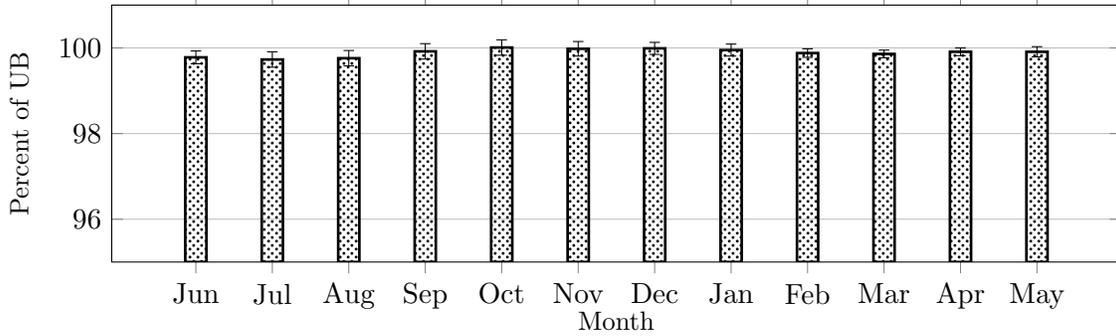
Figure 7: Estimated MLSM-based lower bounds (percentages of UB).

We estimated the MLSM value function approximation using 10,000 regression sample paths of the array of forward curves ($H = 10,000$). We used this value function approximation and 30,000 simulation sample paths of the array of forward curves ($L = 30,000$) to estimate the greedy lower bound and dual upper bound, presented in §5 and Online Appendix C, respectively, on the time T_0 combined market value of the storage and transport assets, $V_0(x_0, \boldsymbol{F}_0)$. Figure 7 reports the lower bound estimates (LBs) as percentages of the dual upper bound estimates (UBs). The error bars in Figure 7 are the standard errors of the LBs, which are at most 0.15% of their respective UBs (the standard errors of the UBs are at most 0.02% of their respective values). The MLSM-based LBs and UBs are essentially tight on all our instances, a finding consistent with the results of Nadarajah et al. (2014a) for the case of natural gas storage with access to a single market.

Our computational setup is a 64 bits PowerEdge R515 with twelve AMD Opteron 4176 2.4GHz processors with 64GB of memory, the Linux Fedora 17 operating system, and the gcc version 4.7.2 20120921 (Red Hat 4.7.2-2) compiler. We use the LAPACK 3.X library with a single processor for ordinary least squares regression and Gurobi 5.0 (Gurobi Optimization 2012) for solving linear programs. Estimating the value function approximations and the MLSM-based bounds takes about 14 minutes per instance. Roughly, 17%, 31%, and 52% of this time is used for estimating a value

function approximation, a lower bound, and an upper bound, respectively.

6.3   Tradeoff between Storage and Transport Trading

To quantify the tradeoff between storage and transport trading we estimate the market value of the transport-only policy, $V_0^{TR}(x_0, \boldsymbol{F}_0)$, and a lower bound on the market value of storage-only policy, $V_0^{ST}(x_0, \boldsymbol{F}_0)$ (see Proposition 3.1). We estimate $V_0^{TR}(x_0, \boldsymbol{F}_0)$ based on solving a linear program for each decision date along a Monte Carlo sample path of the simulation of the array of forward curves. For the estimation of a lower bound on $V_0^{ST}(x_0, \boldsymbol{F}_0)$ we use a simplified version of our MLSM method that ignores transport trades. There is considerable tradeoff between storage and transport trading on our instances. Specifically, our estimates of the average values, across all our instances, of $V_0^{TR}(x_0, \boldsymbol{F}_0)$ and of the lower bound on $V_0^{ST}(x_0, \boldsymbol{F}_0)$ are 75% and 55% of the LB average, respectively, which combined amount to 130% of the LB average.

   To assess the difficulty of managing this tradeoff we consider two sequential policies that prioritize storage over transport trading and vice versa. At a given stage and state, the policy that prioritizes storage over transport first optimizes the storage decisions and then optimizes the transport decisions given the residual transport capacity; the second policy reverses the order of this prioritization. We estimate lower bounds on the market values of these policies by using obvious modifications of our MLSM method. Management of the observed tradeoff between storage and transport trading is difficult: On average, the estimated lower bounds on the optimal values of the policies that respectively prioritize storage over transport and transport over storage are 90% and 78% of LB, where the better performance of the first policy is due to its (intuitively) superior ability to assess the future consequences of current decisions. This finding suggests that there is substantial benefit in practice from jointly optimizing natural gas storage and transport trading decisions in a network setting.

6.4   Extended Rolling Intrinsic Policy and its Simplified Versions

When managing storage in a single market setting, reoptimizing (rolling) at every stage and observed state a policy that computes the intrinsic value of storage, that is, the market value of storage assuming that there is no uncertainty in the evolution of the forward curve, is a near optimal and common approach among practitioners to capture the value of storage, known as the rolling intrinsic policy (Gray and Khandelwal 2004, Breslin et al. 2009, Lai et al. 2010, Secomandi 2010b,

2014, Secomandi et al. 2014, Secomandi and Wang 2012). In our setting, the intrinsic value of the storage and transport assets at the initial stage and state is the optimal value of the deterministic version of the MDP (7) that only depends on the time T_0 array of forward curves $\boldsymbol{F}_0$. In other words, it is the value of the optimal operating policy that relies only on this information. We denote this intrinsic value by $V_0^I(y_0; \boldsymbol{F}_0)$, which can be computed by solving the following deterministic version of SDP (16):

$$V_i^I(y_i; \boldsymbol{F}_0) = \max_{y_{i+1} \in \mathcal{Z}'(y_i)} \bar{r}(y_i - y_{i+1}, \boldsymbol{F}_{0,i}) + \delta V_{i+1}^I(y_{i+1}; \boldsymbol{F}_0), \tag{22}$$

$\forall (i, y_i) \in \mathcal{I} \times \mathcal{Y}'$, with boundary conditions $V_N^I(y_N; \boldsymbol{F}_0) := 0$, $\forall y_N \in \mathcal{Y}'$.

The extrinsic value of the storage and transport assets is the part of the value of these assets attributable to price uncertainty, that is, the difference $V_0(y_0, \boldsymbol{F}_0) - V_0^I(y_0; \boldsymbol{F}_0)$, and measures the incremental value that can be gained by adapting the assets' operating policy to the uncertain evolution of the array of forward curves. We estimate the extrinsic value of these assets on our instances by subtracting $V_0^I(y_0; F_0)$ from LB. The estimated extrinsic values are substantial, ranging between 7.17% and 14.47% of the LBs across our instances (with standard errors smaller than 0.2% of their respective LBs) and averaging to 10.41%. Employing a dynamic policy is thus meaningful on these instances.

With a single storage asset, the rolling intrinsic policy aims at heuristically capturing the total (intrinsic plus extrinsic) value of storage by sequentially updating the intrinsic policy at every stage and observed state. We extend the rolling intrinsic policy to our setting by reformulating and reoptimizing the deterministic dynamic program (22) in every stage and state. That is, at stage i and state $(y_i, \boldsymbol{F}_i)$ this dynamic program depends on $\boldsymbol{F}_i$ rather than $\boldsymbol{F}_0$. This policy implements only the storage and transport decisions pertaining to this stage and state. It then reoptimizes the dynamic program (22) at the stage $i + 1$ state obtained from performing these decisions and observing the new array of forward curves $\boldsymbol{F}_{i+1}$, implements only the storage and transport decisions corresponding to this stage and state, and repeats this process until the end of the horizon is reached. A version of this extended rolling intrinsic policy is part of the commercial software StoragePLUS (FEA 2013).

We estimate the market value of the extended rolling intrinsic policy via Monte Carlo simulation

of the array of forward curves and of the storage inventory level induced by this policy. On our instances this policy performs near optimally, its estimated market value being within 1 standard error of UB on every instance. This finding provides some numerical support for using the extended rolling intrinsic policy in practice.
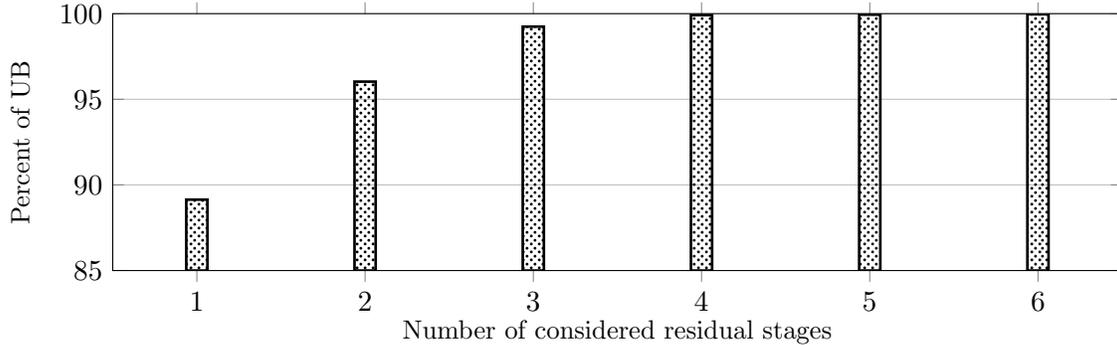
Figure 8: Average estimated lower bounds corresponding to the modified extended rolling intrinsic policies (percentages of UB).

Despite the near optimality of the extended rolling intrinsic policy, estimating the value of this policy is computationally intensive, requiring at least 2 orders of magnitude more CPU time compared to the estimation of (the also near optimal) LB. We thus investigate the performance of simplified extended rolling intrinsic policies based on reoptimizing simpler, and hence less onerous to solve, versions of the intrinsic model (22) formulated with fewer residual stages than there actually are at a given stage. For example, if the number of residual stages is six, then each version of model (22) that we solve includes at most six stages (fewer stages are included when the actual number of remaining stages is less than six).

Figure 8 plots as percentages of the UBs the average estimated lower bounds for the simplified extended rolling intrinsic policies corresponding to different choices of the number of considered residual stages in each optimization. This simplified approach achieves near optimal performance when using as few as 4 stages and leads to computational savings of about 1 order of magnitude relative to the extended rolling intrinsic policy. However, it remains substantially slower than our MLSM approach for the purpose of estimating a lower bound on the combined market value of the storage and transport assets.

6.5 Managing Storage Jointly with Transport Assets Versus as a Single Asset

Storage and transport trades compete for pipeline capacity. To understand the effect of this competition on the management of the storage asset on our instances, we compare three performance metrics when using near-optimal MLSM-based policies that manage the storage asset both jointly with transport assets and as a single asset. Specifically, in both cases we measure the average number of natural gas units flowing through storage (the flow rate, in MMBtu/month) and the average inventory (in MMBtu), and then use Little's law to compute the average amount of time a unit spends in storage (the average flow time, in months). We report the values of these metrics averaged across our instances.

Intuitively, it is expected that the flow rate and average flow time when storage and transport assets are jointly managed will be smaller and larger, respectively, than their corresponding values when managing storage as a single asset. These differences are substantial on our instances: The storage flow rate and average flow time with jointly managed storage and transport assets are 63% and 153% of their corresponding values, 0.2213 MMBtu/month and 2.2 months, respectively, when managing a single storage asset (recall that the storage asset space is 1 MMBtu in our instances). Less obvious is that the average inventory when storage and transport assets are jointly managed is 95% of the average inventory when managing storage alone (0.4814 MMBtu). This reduction in average inventory can be attributed to the flow rate reduction being stronger than the increase in the average flow time.

As expected, competition between storage and transport trades also reduces the average storage margin. However, this reduction is moderate. Specifically, the average storage margin with jointly managed storage and transport assets is 96% of the average storage margin when managing storage alone ($2.579 /MMBtu). In particular, this decrease results from a slight decrease of the average selling price in the former case compared to the latter case, while the average purchase price remains largely unchanged on our instances.

7 Conclusions

The operations of merchant energy trading in wholesale markets across different locations and times can be represented as a network where storage and transport trades compete for the limited capacity

of storage and transport assets. We investigate the tradeoff between storage and transport trades for the realistic case of a network with a single storage asset and multiple transport assets, which we model as an MDP. Because computing an optimal policy of this MDP is intractable, we use our structural analysis of this model to modify an LSM-based method to compute a heuristic policy and estimate lower and upper bounds on the market value of an optimal policy.

On a realistic natural gas application, we observe a substantial and difficult to manage tradeoff between storage and transport trading, which our LSM-based heuristic policy, being near optimal, manages effectively. Although equally effective, a practice-based method, based on sequentially reoptimizing a deterministic model, is, even after simplification, considerably more computationally intensive than our approach. We also highlight the operational differences between managing storage jointly with transport assets versus as a single asset. Our research provides us with an improved understanding of the tradeoff between natural gas storage and transport trading on a network, and offers natural gas merchants both a more efficient method to near optimally manage storage in a network than is possible with modifications of a method currently available in practice and a way to assess the suboptimality of heuristics.

Beyond natural gas, our research has relevance for managing the merchant trading operations of other energy sources, such as coal, electricity, and oil and petroleum products, natural resources, such as water and timber, and other storable commodities, such as agricultural products and metals. Further research could focus on the management of the merchant operations of energy trading for the case of a network that includes multiple storage assets, in addition to multiple transport assets (see Löhndorf et al. 2013 and Salas and Powell 2014 for recent related work).

Acknowledgments

References

Arvesen, Ø., V. Medbø, S. E. Fleten, A. Tomasgard, S. Westgaard. 2013. Linepack storage valuation under price uncertainty. *Energy* **52**(1) 155–164.

Bannister, C. H., R. J. Kaye. 1991. A rapid method for optimization of linear systems with storage. *Operations Research* **39**(2) 220–232.

Bertsekas, D. P. 2012. *Dynamic Programming and Optimal Control, Volume II*. Athena Scientific, Belmont, Massachusetts, USA.

Björk, T. 2004. *Arbitrage Theory in Continuous Time*. Oxford University Press, Oxford, Oxfordshire, UK.

Blanco, C., D. Soronow, P. Stefiszyn. 2002. Multi-factor models for forward curve analysis: An introduction to principal component analysis. *Commodities Now* **June** 76–78.

Boogert, A., C. De Jong. 2011/12. Gas storage valuation using a multifactor price process. *The Journal of Energy Markets* **4**(4) 29–52.

Boogert, A., D. Mazières. 2013. A radial basis function approach to gas storage valuation. *The Journal of Energy Markets* **6**(2) 19–50.

Boyabatli, O. 2011. Supply management in multi-product firms with fixed proportions technology. Working paper, Singapore Management Univ.

Boyabatli, O., P. R. Kleindorfer, S. R. Koontz. 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science* **57**(10) 1771–1787.

Breslin, J., L. Clewlow, T. Elbert, C. Kwok, C. Strickland, D. van der Zee. 2009. Gas storage: Rolling instrinsic valuation. *Energy Risk* **January** 61–65.

Brown, D. B., J. E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4) 1–17.

Charnes, A., J. Drèze, M. Miller. 1966. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* **34**(2) 307–330.

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

Cortazar, G., E. S. Schwartz. 1994. The valuation of commodity contingent claims. *The Journal of Derivatives* **1**(4) 27–39.

Deng, S., B. Johnson, A. Sogomonian. 2001. Exotic electricity options and the valuation of electricity generation and transmission assets. *Decision Support Systems* **30**(3) 383–392.

Devalkar, S. K., R. Anupindi, A. Sinha. 2011. Integrated optimization of procurement, processing, and trade of commodities. *Operations Research* **59**(6) 1369–1381.

Dixit, A. K., R. S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, New Jersey, USA.

EIA. 2013. Anual Energy Outlook 2013. Tech. rep., Energy Information Agency (EIA).

Eydeland, A., K. Wolyniec. 2003. *Energy and Power Risk Management*. John Wiley and Sons Inc., Hoboken, NJ, USA.

FEA. 2013. StoragePLUS features. Tech. rep., Financial Engineering Associates (FEA), `http://www.fea.com/products/energy/storage_plus/features.html`.

Frestad, D. 2008. Common and unique factors influencing daily swap returns in the Nordic electricity market, 1997-2005. *Energy Economics* **30**(3) 1081–1097.

Gabriel, S. A., S. Kiet, J. Zhuang. 2005. A mixed complementarity-based equilibrium model of natural gas markets. *Operations Research* **53**(5) 799–818.

Geman, H. 2005. *Commodities and Commodity Derivatives: Modelling and Pricing for Agriculturals, Metals, and Energy*. Wiley, Chichester, UK.

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

Gray, J., P. Khandelwal. 2004. Realistic gas storage models II: Trading strategies. *Commodities Now* **September** 1–5.

Gurobi Optimization, Inc. 2012. Gurobi optimizer reference manual version 5.0. Houston, Texas: Gurobi Optimization.

INGAA. 2009. Natural gas pipeline and storage infrastructure projections through 2030. Tech. rep., Interstate Natural Gas Association of America (INGAA).

Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50**(7) 909–921.

Kim, J. H., W. B. Powell. 2011. Optimal energy commitments with storage and intermittent supply. *Operations Research* **59**(6) 1347–1360.

Kleindorfer, P. R., D. J. Wu. 2003. Integrating long-and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science* **49**(11) 1597–1615.

Lai, G., F. Margot, N. Secomandi. 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58**(3) 564–582.

Lai, G., M. X. Wang, S. Kekre, A. Scheller-Wolf, N. Secomandi. 2011. Valuation of storage at a liquefied natural gas terminal. *Operations Research* **59**(3) 602–616.

Löhndorf, N., S. Minner. 2010. Optimal day-ahead trading and storage of renewable energies: An approximate dynamic programming approach. *Energy Systems* **1**(1) 61–77.

Löhndorf, N., D. Wozabal, S. Minner. 2013. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* **61**(4) 810–823.

Longstaff, F. A., E. S. Schwartz. 2001. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* **14**(1) 113–147.

Luenberger, D. G. 1998. *Investment Science*. Oxford University Press, New York, New York, USA.

Maragos, S. 2002. Valuation of the operational flexibility of natural gas storage reservoirs. E. Ronn, ed., *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 431–456.

Markland, R. E. 1975. Analyzing multi-commodity distribution networks having milling-in-transit features. *Management Science* **21**(12) 1405–1416.

Markland, R. E., R. J. Newett. 1976. Production-distribution planning in a large scale commodity processing network. *Decision Sciences* **7**(4) 579–594.

Nadarajah, S., F. Margot, N. Secomandi. 2014a. Comparison of least squares Monte Carlo methods with applications to energy real options. Working paper, Carnegie Mellon Univ.

Nadarajah, S., F. Margot, N. Secomandi. 2014b. Relaxations of approximate linear programs for the real option management of commodity storage. Working paper, Carnegie Mellon Univ.

NAESB. 2002. Base contract for retail sale and purchase of natural gas or electricity, section 3. Tech. rep., North American Energy Standards Board (NAESB).

Nascimento, J., W. Powell. 2013. An optimal approximate dynamic programming algorithm for concave, scalar storage problems with vector-valued controls. *IEEE Transactions on Automatic Control* **58**(12) 2995–3010.

Pipeline Knowledge & Development. 2010. The interstate natural gas transmission system: Scale, physical complexity and business model. Tech. rep., The Interstate Natural Gas Association of America (INGAA). Http://www.ingaa.org/11885/Reports/10724.aspx.

Powell, W. B. 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality, 2nd Edition*. John Wiley & Sons, Hoboken, New Jersey, USA.

Rogers, L. C. G. 2002. Monte Carlo valuation of American options. *Mathematical Finance* **12**(3) 271–286.

Rømo, F., A. Tomasgard, L. Hellemo, M. Fodstad, B. H. Eidesen, B. Pedersen. 2009. Optimizing the Norwegian natural gas production and transport. *Interfaces* **39**(1) 46–56.

Salas, D., W. B. Powell. 2014. Benchmarking a scalable approximation dynamic programming algorithm for stochastic control of multidimensional energy storage problems. Working paper, Princeton Univ.

Schwartz, E., J. E. Smith. 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46**(7) 893–911.

Scott, W. R., W. B. Powell, S. Moazeni. 2014. Least squares policy iteration with instrumental variables vs. direct policy search: Comparison against optimal benchmarks using energy storage. Working paper, Princeton Univ.

Secomandi, N. 2010a. On the pricing of natural gas pipeline capacity. *Manufacturing & Service Operations Management* **12**(3) 393–408.

Secomandi, N. 2010b. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 1090–1049.

Secomandi, N. 2014. Analysis and enhancement of practice-based policies for the real option management of commodity storage assets. Working paper, Carnegie Mellon Univ.

Secomandi, N., G. Lai, F. Margot, A. Scheller-Wolf, D. Seppi. 2014. Merchant commodity storage and term structure model error. Working paper, Carnegie Mellon Univ.

Secomandi, N., D. Seppi. 2014. Real Options and Merchant Operations of Energy and Other Commodities. *Foundations and Trends in Technology, Information and Operations Management* **6**(3-4) 161–331.

Secomandi, N., M. X. Wang. 2012. A computational approach to the real option management of network contracts for natural gas pipeline transport capacity. *Manufacturing & Service Operations Management* **14**(3) 441–454.

Smith, J. E., K. F. McCardle. 1998. Valuing oil properties: Integrating option pricing and decision analysis approaches. *Operations Research* **46**(2) 198–217.

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

Smith, R. 2013. Can gas undo nuclear power? *Wall Street Journal* .

Sturm, J. F. 1997. *Trading Natural Gas: A Nontechnical Guide*. PennWell Publishing Company, Tulsa, OK.

Suenaga, H., A. Smith, J. Williams. 2008. Volatility dynamics of NYMEX natural gas futures prices. *Journal of Futures Markets* **28**(5) 438–463.

Thompson, M. 2012. Natural gas storage valuation, optimization, market and credit risk management. Working paper, Queens Univ.

Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, New Jersey.

Tsitsiklis, J. N., B. Van Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* **12**(4) 694–703.

Wu, O. Q., D. D. Wang, Z. Qin. 2012. Seasonal energy storage operations with limited flexibility: The price-adjusted rolling intrinsic policy. *Manufacturing & Service Operations Management* **14**(3) 455–471.

Zhou, Y., A. Scheller-Wolf, N. Secomandi, S. Smith. 2013. Managing wind-based electricity generation in the presence of storage and transmission capacity. Working paper, Carnegie Mellon Univ.

Zhou, Y., A. Scheller-Wolf, N. Secomandi, S. Smith. 2014. Merchant management of electricity surpluses: Storage vs. disposal. Working paper, Carnegie Mellon Univ.

Online Appendix

A Pipeline Figures

Figures 9-10 that supplement the discussion in §2.
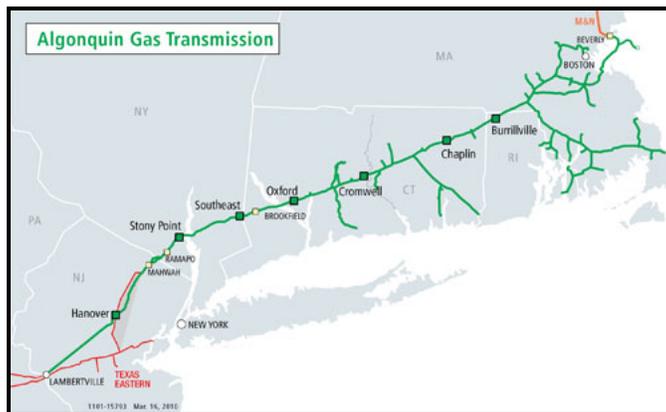
Figure 9: Interconnect stations between the TETCO and AGT pipeline systems (Source: Spectra Energy website).

Figure 10: The TRANSCO pipeline system (Source: Rextag Strategies website).

Figure 11: The AGT pipeline system (Source: Rextag Strategies website).

B   Proofs

The section includes the proofs of the results stated in §3 and §4.2.

Proof of Proposition 3.1. Let π^* be an optimal policy to (7) and $\pi^{ST,*}$ and $\pi^{TR,*}$ be the storage and transport policies, respectively, that make up π^*. The required inequality follows because $\pi^{ST,*}$ and $\pi^{TR,*}$ are feasible but not necessarily optimal policies to the two model version of (7) specified with the restrictions $\pi \in \Pi^{TR}$ and $\pi \in \Pi^{ST}$, respectively. □

Proof of Lemma 4.1. To prove the claimed characterization we require the finiteness of the value and continuation functions of SDP (14). It is obvious that $V_i(y_i, \boldsymbol{F}_i) \geq 0 > -\infty$, which implies that $W_i(y_{i+1}, \boldsymbol{F}_i) > -\infty$. Further, $V_i(y_i, \boldsymbol{F}_i) \leq \sum_{i' \in \mathcal{I}_i} \sum_{m \in \mathcal{M}} C^{D,m} s_{i'}^m$. Using this inequality we have

$$W_i(y_{i+1}, \boldsymbol{F}_i) \equiv \delta \mathbb{E}\left[V_{i+1}\left(y_{i+1}, \boldsymbol{F}_{i+1}\right) | \boldsymbol{F}_i\right]$$

$$\leq \delta \mathbb{E}\left[\sum_{i' \in \mathcal{I}_i} \sum_{m \in \mathcal{M}} C^{D,m} s_{i'}^m | \boldsymbol{F}_i\right]$$

$$= \delta \sum_{i' \in \mathcal{I}_i} \sum_{m \in \mathcal{M}} C^{D,m} \mathbb{E}\left[s_{i'}^m | \boldsymbol{F}_i\right]$$

$$= \delta \sum_{i' \in \mathcal{I}_i} \sum_{m \in \mathcal{M}} C^{D,m} F_{i,i'}^m$$

$$< \infty,$$

where the second equality follows from $s_{i''}^m \equiv F_{i'',i'}^m$ and the martingale property of futures prices (Shreve 2004, page 244). Thus, the value and continuation functions of SDP (14) are finite.

We now proceed by induction to prove the claimed result. At stage $N-1$, for a given $\boldsymbol{F}_i$, we have

$$V_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = \max_{y_N \in \mathcal{Z}(y_{N-1})} \bar{r}(y_{N-1} - y_N, \boldsymbol{s}_{N-1}).$$

Standard linear programming results (Bertsimas and Tsitsiklis 1997, Ch. 5) imply that the function $\bar{r}(y_{N-1} - y_N, \boldsymbol{s}_{N-1})$ is concave in the pair (y_{N-1}, y_N), which belongs to a convex (polyhedral) feasible set. Further, the interval $\mathcal{Z}(y_{N-1})$ is nonempty for each given $y_{N-1} \in \mathcal{Y}$. The concavity of $V_{N-1}(\cdot, \boldsymbol{F}_{N-1})$ follows from Proposition B-4 in Heyman and Sobel (2003). The continuation function at stage $N-1$ is zero by definition and is therefore concave.

Make the induction hypothesis that the value and continuation functions are concave in their first arguments also for stages $i+1, i+2, \ldots, N-2$. We proceed to prove the claim at stage i. From the finiteness of the continuation function in every stage and the induction hypothesis, it is easy to verify that the continuation function is concave in its first argument at stage i. This property and Part (a) of this lemma imply the concavity of $\bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i) + W_i(y_{i+1}, \boldsymbol{F}_i)$ in the pair (y_i, y_{i+1}), which belongs to a convex (polyhedral) feasible set. Since $\mathcal{Z}(y_i)$ is nonempty for each given $y_i \in \mathcal{Y}$, the concavity of $V_i(\cdot, \boldsymbol{F}_i)$ follows from Proposition B-4 in Heyman and Sobel (2003). The claimed concavity of the value and continuation functions at all stages for a given array of forward curves follows from the principle of mathematical induction. $\square$

Proof of Lemma 4.3. The piecewise linear concavity of $\bar{r}(\cdot, \boldsymbol{s})$ follows from standard linear programming results (Bertsimas and Tsitsiklis 1997, Ch. 5). We proceed to show that the slope of $\bar{r}(\cdot, \boldsymbol{s})$ changes at integer multiples of $\bar{G}$. Our proof relies on reformulating the maximum profit flow problem (12), defined over the set of trade paths, as a maximum profit network flow problem on an *edge* network.

We begin by describing the edge network construction. Figure 12 illustrates the edge network assuming a nonnegative storage action, $a \leq 0$. An analogous network exists for the case $a > 0$.
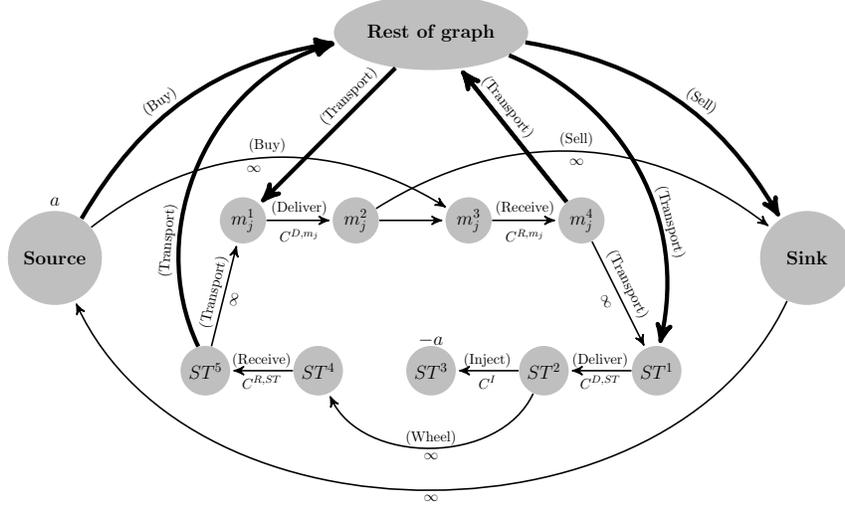
Figure 12: Edge network formulation for the feasible set of (12) when $a \leq 0$.

Recall that $C^{R,m}$ and $C^{D,m}$ are the receipt and delivery capacities of node m, respectively. The nodes in this figure are (i) a dummy source node and a dummy sink node; (ii) nodes m_j^1-m_j^4 for modeling market m_j; (iii) nodes ST^1-ST^5 for modeling storage; and (iv) the node labeled "rest of graph" that is an aggregation of nodes and edges not represented explicitly in this figure. Thin edges denote actual edges in the network with their capacities as edge labels, while thick edges denote a collection of edges to or from the node labeled "rest of the graph". The supply of the source node is the storage action a. The demand of node ST^3 is $-a$. All other nodes have zero demand. It can be verified that each path in this network corresponds to the path of a trade in set $\mathcal{J}$. To help verify this assertion, the labels within parentheses on the edges of this figure indicate the edge modeling purpose. For example, the label "(Buy)" on the edge (Source, m_j^3) indicates that this edge models the purchase of energy from market m_j. Thus, each term that makes up the cash flow of a trade can be represented as a profit on one of the edges in this network.

We introduce some generic notation to formulate the stated maximum profit flow problem on this edge network. Part of this notation uses with a different meaning notation used elsewhere in this paper, but this confined reuse of notation should not give rise to any confusion. The pair $(\mathcal{N}, \mathcal{E})$ includes the node and edge sets of this network. Let the capacity and profit on edge $e = (u, v)$ be C_e and c_e, respectively (c is the column vector of all edge profits). Denote the demand on node v by d_v. By construction, we have $d_{source} = a$, $d_{ST^3} = -a$, and all other node demands equal to zero.

We use w_e to represent the flow variable associated with edge e (w is the column vector of all these flow variables). The maximum profit flow problem is

$$\max_{w} c^{\top} \cdot w \tag{23}$$

$$\sum_{e \in \mathcal{E}: e=(\cdot,v)} w_e - \sum_{e \in \mathcal{E}: e=(v,\cdot)} w_e = 0, \forall v \in \mathcal{N} \setminus \{\text{Source}, ST^3\}, \tag{24}$$

$$w_{(Sink,Source)} - \sum_{e \in \mathcal{E}: e=(Source,\cdot)} w_e = a, \tag{25}$$

$$w_{(ST_2,ST_3)} = -a, \tag{26}$$

$$0 \leq w_e \leq C_e, \forall e \in \mathcal{E}. \tag{27}$$

The claimed equivalence between (23)-(27) and (12) holds because an optimal solution of (23)-(27) can be decomposed into amounts corresponding to trades in set $\mathcal{J}$ by removing the edge from the sink to the source if $a \neq 0$ (all cycles that do not include the edge (Sink, Source) have a negative profit).

Make the change of variable $\hat{w}_e := w_e / \bar{G}$ in (23)-(27) ($\hat{w}$ is the column vector of scaled flow variables on all edges). This change yields the linear program

$$\max_{\hat{w}} \bar{G}(c^{\top} \cdot \hat{w}) \tag{28}$$

$$\sum_{e \in \mathcal{E}: e=(\cdot,v)} \hat{w}_e - \sum_{e \in \mathcal{E}: e=(v,\cdot)} \hat{w}_e = 0, \forall v \in \mathcal{N} \setminus \{\text{Source}, ST^3\}, \tag{29}$$

$$\hat{w}_{(Sink,Source)} - \sum_{e \in \mathcal{E}: e=(Source,\cdot)} \hat{w}_e = a/\bar{G}, \tag{30}$$

$$\hat{w}_{(ST_2,ST_3)} = -a/\bar{G}, \tag{31}$$

$$0 \leq \hat{w}_e \leq C_e/\bar{G}, \forall e \in \mathcal{E}. \tag{32}$$

From the definition of $\bar{G}$ it follows that the scaled capacity $C_e/\bar{G}$ is integer for all the edges that have finite capacity. The integrality of the scaled capacities and the unimodularity of the constraint matrices of network flow problems (see Theorems 11.11 and 11.12 in Ahuja et al. 1993) imply that the optimal solution to the linear program (28)-(32) with $a = 0$ is integer. Let this optimal solution be $\hat{w}^*(0)$. Recall the definition of a^I on Page 14. Suppose that we increase the injection amount

from 0 to ϵ such that $0 < \epsilon < \bar{G}$ and $\epsilon \leq a^I$, that is, $a = -\epsilon$. This corresponds to increasing the injection amount by less than 1 unit in (28)-(32). Note that $\hat{w}^*(0)$ defines a pseudo flow (see page 320 of Ahuja et al. 1993) for the problem with $a = -\epsilon$, that is, this solution violates mass balance by ϵ only at the nodes Source and ST^3. Since $\epsilon \leq a^I$ there exists a shortest path from the node Source to the node ST^3 in the residual network (see §9.1 in Ahuja et al. 1993). Augmenting by ϵ the flow along this path yields an optimal solution to (28)-(32) by Lemma 9.12 in Ahuja et al. (1993). This shortest path has capacity at least 1 since $\hat{w}^*(0)$ is an integral flow and all edge capacities in (28)-(32) are integral. Thus, when changing $a/\bar{G}$ between 0 and -1 an optimal solution to the resulting problem can be found by augmenting by the same amount the flow along the same shortest path. Consequently, the optimal solution value of the problem (28)-(32) is linear for values of $a/\bar{G}$ between 0 and -1, which implies that the optimal solution to (23)-(27) is linear for $a \in [0, -\bar{G}]$. This argument can be repeated to prove an analogous result when the injection amount ϵ is between any two consecutive integers η and $\eta + 1$ such that $(\eta + 1)\bar{G} \leq a^I$. These arguments prove the claimed result for the injection case. Symmetric arguments can be used to show the claimed result for withdrawals. □

Proof of Proposition 4.4. By induction.

Stage $N - 1$.

By Lemma 4.3, $\bar{r}_{N-1}(\cdot, \boldsymbol{s}_{N-1})$ is piecewise linear concave with slope changes at integer multiples of $\bar{G}$, and thus at integer multiples of G as well. Thus, $\max_{a \in \mathbb{R}} \bar{r}_{N-1}(a, \boldsymbol{s}_{N-1})$ has a maximizer a_{N-1}^* that is an integer multiple of G, where we suppress the dependence of this maximizer on $\boldsymbol{s}_{N-1}$. This maximizer is also finite, for every given $\boldsymbol{s}_{N-1}$, because, as discussed in §4.1, the linear program (12) is infeasible when $a \notin [-a^I, a^W]$. Notice that $a_{N-1}^* \geq 0$ since injecting ($a < 0$) incurs an additional cost compared to doing nothing ($a = 0$). If $a_{N-1}^* = 0$, then $b_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = y_{N-1}$ for all $y_{N-1} \in [0, \bar{y}]$, which is a linear function in y_{N-1} with slope equal to 1. Suppose that $a_{N-1}^* > 0$, that is, withdrawal is optimal. Then, it holds that $b_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = 0$ for all $y_{N-1} \in [0, \min\{a_{N-1}^*, \bar{y}\}]$ and $b_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = y_{N-1} - a_{N-1}^*$ for all $y_{N-1} \in (\min\{a_{N-1}^*, \bar{y}\}, \bar{y}]$. This target function is piecewise linear in y_{N-1} with possible slopes equal to 0 and 1. Thus, we have established the claimed structure of the target function at stage $N - 1$. This structure implies
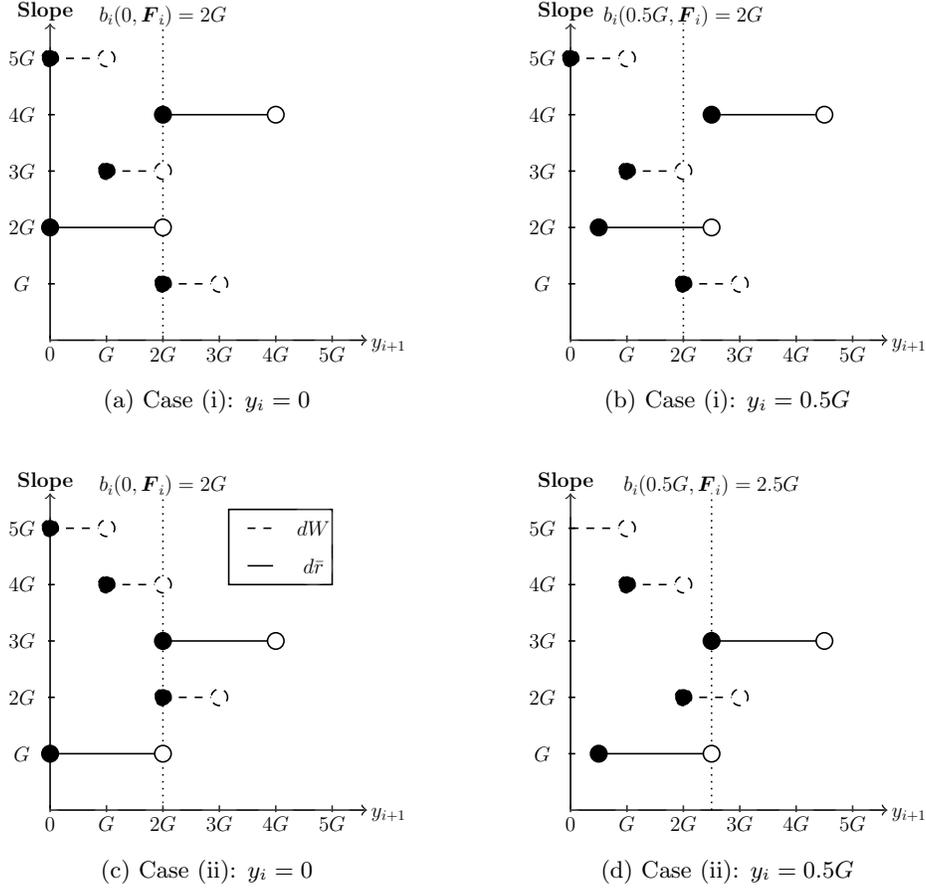
Figure 13: Conceptual illustration of cases (i) and (ii) in the proof of Part (b) of Proposition 4.4.

that $\underline{b}_{N-1}(\boldsymbol{F}_{N-1}) = 0$, while $\bar{b}_{N-1}(\boldsymbol{F}_{N-1}) = \bar{y}$ if $a^*_{N-1} = 0$ and $\bar{b}_{N-1}(\boldsymbol{F}_{N-1}) = 0$ if $a^*_{N-1} > 0$, which is consistent with our claimed partitioning of the feasible inventory set $\mathcal{Y}$. Finally, the target function structure also implies that the difference $y_{N-1} - b_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1})$ is not increasing in y_{N-1} and the values returned by this function lie in the same region of the inventory partition that y_{N-1} belongs to.

By definition, $W_{N-1}(\cdot, \boldsymbol{F}_{N-1})$ is constant at zero, and thus trivially satisfies the claimed property. By Lemma 4.1, $V_{N-1}(\cdot, \boldsymbol{F}_{N-1})$ is concave. We proceed to prove that this function is piecewise linear concave with breakpoints at integer multiples of G. Recall from above that $a^*_{N-1} \equiv \text{argmax}_{a \in \mathbb{R}} \bar{r}_{N-1}(\cdot, \boldsymbol{s}_{N-1})$. Define $\bar{a} := \min\{C^W, a^W, a^*_{N-1}, \bar{y}\}$, which is an integer multiple of G. For $y_{N-1} \in [0, \bar{a})$ it holds that $V_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = \bar{r}_{N-1}(y_{N-1}, \boldsymbol{s}_{N-1})$, and for $y_{N-1} \in [\bar{a}, \bar{y}]$ we have $V_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1}) = \bar{r}_{N-1}(\bar{a}, \boldsymbol{s}_{N-1})$. The function $V_{N-1}(y_{N-1}, \boldsymbol{F}_{N-1})$ inherits the slope of $\bar{r}_{N-1}(y_{N-1}, \boldsymbol{s}_{N-1})$ for $y_{N-1} \in [0, \bar{a}]$ and has a slope of zero for $y_{N-1} \in [\bar{a}, \bar{y}]$. Because the function

$\bar{r}_{N-1}(\cdot, \boldsymbol{s}_{N-1})$ is piecewise linear concave with slope changes slope at integer multiples of G, so is $V_{N-1}(\cdot, \boldsymbol{F}_{N-1})$.

Induction hypothesis. Suppose that the value function $V_i(\cdot, \boldsymbol{F}_i)$ is piecewise linear concave with slope changes at integer multiples of G for stages $i' = i+1, \ldots, N-2$.

Stage i. Because $V_{i+1}(y_{i+1}, \boldsymbol{F}_{i+1})$ is finite (see the proof of Lemma 4.1), for a fixed $\boldsymbol{F}_{i+1}$, the induction hypothesis implies that $W_i(y_{i+1}, \boldsymbol{F}_i) \equiv \delta \mathbb{E}[V_{i+1}(y_{i+1}, \boldsymbol{F}_{i+1})|\boldsymbol{F}_i]$ is a piecewise linear function with slope changes at integer multiples of G.

For a given y_i, note that $d\bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i)/dy_{i+1}|_{y_{i+1}=y'_{i+1}} = -d\bar{r}(a, \boldsymbol{s}_i)/da|_{a=y_i-y'_{i+1}}$. Moreover, it holds that $d\bar{r}(a, \boldsymbol{s}_i)/da|_{a=y_i-y'_{i+1}}$ is a nondecreasing function of y'_{i+1} because the function $\bar{r}(\cdot, \boldsymbol{s}_i)$ is concave. Since $W_i(\cdot, \boldsymbol{F}_i)$ and $\bar{r}(\cdot, \boldsymbol{s}_i)$ are piecewise linear concave functions, we can state the optimality condition that determines $b_i(y_i, \boldsymbol{F}_i)$ as follows: $b_i(y_i, \boldsymbol{F}_i)$ is the smallest $y'_{i+1} \in \mathcal{Y}$ such that $d\bar{r}(a, \boldsymbol{s}_i)/da|_{a=y_i-y'_{i+1}} \geq dW_i(y_{i+1}, \boldsymbol{F}_i)/dy_{i+1}|_{y_{i+1}=y'_{i+1}}$.

Before using this optimality condition to formally prove the structure in y_i of the target function, we provide the intuition behind our proof. Observe that as y_i increases the slope of the continuation function $W_i(\cdot, \boldsymbol{F}_i)$ does not change but the slope of the reward function $\bar{r}_i(\cdot, \boldsymbol{s}_i)$ weakly decreases. In other words, $dW_i(\cdot, \boldsymbol{F}_i)/dy_{i+1}|_{y_{i+1}=y'_{i+1}}$ is a nonincreasing step function with changes at integer multiples of G that does not depend on y_i, whereas $d\bar{r}(a, \boldsymbol{s}_i)/da|_{a=y_i-y'_{i+1}}$ is a nondecreasing step function with changes at integer multiples of G that shifts to the right by the same proportion by which y_i is increased. Figure 13 provides illustrative examples of this property. Our proof of the basestock target function structure relies on two possible types of optima. The first type of optimum, illustrated in panels (a) and (b) of Figure 13 occurs when the slope of the reward function is "bracketed" above and below by the slopes of the continuation function as y_i is increased starting from a value that is an integer multiple of G. In this case, the target function is a constant until the reward function slope is no longer bracketed. The second type of optimum, illustrated in panels (c) and (d) of Figure 13 occurs when the continuation function slope is bracketed between the reward function slopes as y_i is increased. In this case, the target function increases proportionately to the increase in y_i until the continuation function slope is no longer bracketed. We now make these intuitive arguments formal.

Let $dW(q)$ be the slope of the function $W_i(\cdot, \boldsymbol{F}_i)$ in the interval $[qG, (q+1)G)$ for $q = 0, 1, \ldots, (\bar{y}/G) - 1$. By Lemma 4.3, $\bar{r}_i(\cdot, \boldsymbol{s}_i)$ changes slope at integer multiples of $\bar{G}$, and thus at integer multiples of G as well. Define $d\bar{r}(q) := d\bar{r}(a, \boldsymbol{s}_i)/da|_{a=(q-1)\bar{G}}$ for $q = -(a^I/G) + 1, -(a^I/G) + 2, \ldots, a^W/G$. We define these slope functions to be right continuous at breakpoints, except for the right boundary point, where these functions are defined to be left continuous. Because the function $\bar{r}_i(-y_{i+1}, \boldsymbol{s}_i) + W_i(y_{i+1}, \boldsymbol{F}_i)$ changes slope in y_{i+1} at integer multiples of G, it holds that $b_i(0, \boldsymbol{F}_i) = \bar{q}G$ for some nonnegative integer $\bar{q} \leq \min\{a^I/G, \bar{y}/G\}$. Suppose that $0 < \bar{q} < \min\{a^I/G, \bar{y}/G\}$ (we discuss the boundary cases later). The optimality condition stated above implies that the following conditions must hold for $\bar{q}G$ to be optimal at $y_i = 0$:

$$d\bar{r}(\lfloor \bar{q}G/\bar{G} \rfloor) \geq dW(\bar{q}), \tag{33}$$

$$d\bar{r}(\lfloor (\bar{q}G - \epsilon)/\bar{G} \rfloor) < dW(\bar{q} - 1), \text{ for some } \epsilon > 0, \tag{34}$$

where $\lfloor \rfloor$ is the floor function. In addition, we could either have: (ai) $d\bar{r}(\lfloor (\bar{q}G - \epsilon)/\bar{G} \rfloor) \geq dW(\bar{q})$ (see Figure 13(a) for an example) or (aii) $d\bar{r}(\lfloor (\bar{q}G - \epsilon)/\bar{G} \rfloor) < dW(\bar{q})$ (see Figure 13(c) for an example). We now characterize an interval $[0, \eta G) \subseteq \mathcal{Y}$ by suitably defining a positive integer η such that the target function has slope 0 or 1 if case (i) or (ii) holds, respectively.

Case (i). Find the smallest positive integer η such that $d\bar{r}(\lfloor (\bar{q} - \eta)G/\bar{G} \rfloor) < dW(\bar{q})$. When such an η does not exist, set η equal to $\bar{q}/G$ (that is, we reach the left boundary of $\mathcal{Y}$). Therefore, for all $y_i \in [0, \eta G)$ it holds that $d\bar{r}(\lfloor (\bar{q}G - y_i)/\bar{G} \rfloor) \geq dW(\bar{q})$. Further, it holds that there exists an $\epsilon' > 0$ such that $d\bar{r}(\lfloor (\bar{q}G - y_i - \epsilon')/\bar{G} \rfloor) < dW(\bar{q} - 1)$. This inequality follows from (34) and the concavity of $\bar{r}_i(\cdot, \boldsymbol{s}_i)$. Thus, we have $b_i(y_i, \boldsymbol{F}_i) = \bar{q}G, \forall y_i \in [0, \eta G]$. Panels (a) and (b) of Figure 13 provide illustrative examples.

Case (ii). Find the smallest positive integer η such that $d\bar{r}(\lfloor (\bar{q}G - \epsilon)/\bar{G} \rfloor) \geq dW(\bar{q} + \eta)$. When such an η does not exist, set η equal to $(\bar{y} - \bar{q}G)/G$ (that is, we reach the right boundary of $\mathcal{Y}$). Therefore, for all $y_i \in [0, \eta G)$, there exists an $\epsilon' > 0$ such that $d\bar{r}(\lfloor (\bar{q}G - \epsilon)/\bar{G}) < dW(\lfloor (\bar{q}G + y_i - \epsilon')/G \rfloor)$. Moreover, by (33) and the concavity of $W_i(\cdot, \boldsymbol{F}_i)$ we have $d\bar{r}(\lfloor \bar{q}G/\bar{G} \rfloor) \geq dW(\lfloor (\bar{q}G + y_i)/G \rfloor)$. Thus, it holds that $b_i(y_i, \boldsymbol{F}_i) = \bar{q}G + (y_i - \bar{q}G), \forall y_i \in [0, \eta G]$. Panels (a) and (b) of Figure 13 provide illustrative examples.

Irrespective of case (i) or (ii), $b_i(\eta G, \boldsymbol{F}_i) = \bar{q}G$. However, at $y_i = \eta G$ the type of optimum changes from the type of optimum at $y_i = 0$, that is, if case (i) is true at $y_i = 0$ then case (ii) holds at $y_i = \eta G$, and if case (ii) is true at $y_i = 0$ then case (i) holds at $y_i = \eta G$. We then repeat the procedure described above to identify a positive integer η' such that the target function either is constant or increases within the interval $[\eta G, \eta' G] \subset \mathcal{Y}$. This process is iterated until we reach the right boundary of $\mathcal{Y}$.

Now we consider the boundary cases. When $\bar{q} = 0$ a proof analogous to the interior case handled above, omitted for brevity, establishes the claimed result. When $\bar{q} = \min\{a^I/G, \bar{y}/G\}$, we have $b_i(y_i, \boldsymbol{F}_i) = \bar{q}G, \forall y_i \in \mathcal{Y}$. We have thus proved the claimed piecewise linearity of the target function at stage i.

We now show the partitioning of the feasible inventory set into the stated inject, do nothing, and withdraw regions. It is obvious that $-b_i(0, \boldsymbol{F}_i) \leq 0$ and $\bar{y} - b_i(\bar{y}, \boldsymbol{F}_i) \geq 0$. Our characterization of $b_i(y_i, \boldsymbol{F}_i)$ as a piecewise linear function of y_i with slopes 0 or 1 implies that (i) $b_i(y_i, \boldsymbol{F}_i)$ is continuous and nondecreasing in y_i, and (ii) $y_i - b_i(y_i, \boldsymbol{F}_i)$ is a nondecreasing function of y_i. The first property implies that the set $\{y_i | y_i = b_i(y_i, \boldsymbol{F}_i)\}$ is a nonempty closed interval. Thus, the functions $\underline{b}_i(\boldsymbol{F}_i)$ and $\bar{b}_i(\boldsymbol{F}_i)$ are well defined. The second property implies that an optimal storage action does not increase in y_i, which proves the partitioning of the inventory interval into the inject, do nothing, and withdraw regions and the target function returns values that lie in the region of this partition that y_i belongs to.

When y_i is within an interval $[qG, (q+1)G]$ the target function is either (i) equal to a constant, which implies that an optimal storage action increases at rate 1 in y_i and the value function $V_i(\cdot, \boldsymbol{F}_i)$ inherits the slope of $\bar{r}_i(\cdot, \boldsymbol{s}_i)$ or (ii) increases with slope 1, the optimal action is a constant, and the value function $V_i(\cdot, \boldsymbol{F}_i)$ inherits the slope of the continuation function $W_i(\cdot, \boldsymbol{F}_i)$. Since both $\bar{r}_i(\cdot, \boldsymbol{s}_i)$ and $W_i(\cdot, \boldsymbol{F}_i)$ are piecewise linear functions with slope changes at multiples of G, the claimed result follows.

The claimed properties hold for all the stages by the principle of mathematical induction. The decision rule (15) tries to change the current storage inventory level to the basestock target function value while accounting for the storage injection and withdrawal capacities. $\square$

C Dual Upper Bound

In this section we discuss the estimation of a dual upper bound on the market value of an optimal policy of our MDP (Brown et al. 2010 and references therein). Such a bound is based on performing hindsight optimizations in which knowledge of future information is penalized using dual penalties. We denote by $u_i(y_{i+1}, \boldsymbol{F}_i, \boldsymbol{F}_{i+1})$ the dual penalty in stage i given y_{i+1}, $\boldsymbol{F}_i$, and $\boldsymbol{F}_{i+1}$. Specifically, these penalties penalize knowledge in stage i of the array of forward curves $\boldsymbol{F}_{i+1}$ and must satisfy the feasibility condition $\mathbb{E}[u_i(y_{i+1}, \boldsymbol{F}_i, \boldsymbol{F}_{i+1})|\boldsymbol{F}_i] \leq 0$ (see Brown et al. 2010 for details). Once feasible dual penalties are specified, we estimate dual upper bounds in Monte Carlo simulation using the same set of L simulation sample paths $\{\boldsymbol{F}_i^l, i \in \mathcal{I} \setminus \{0\}, l = 1, 2, \ldots, L\}$ employed for lower bound estimation. A point estimate $U_0^l(y_0)$ of an upper bound on $V_0(y_0, \boldsymbol{F}_0)$ can be computed based on the l-th sample path of the arrays of forward curves by solving the following dynamic program, $\forall(i, y_i) \in \mathcal{I} \times \mathcal{Y}'$:

$$U_i^l(y_i) = \max_{y_{i+1} \in \mathcal{Z}_i'(y_i)} \bar{r}(y_i - y_{i+1}, \boldsymbol{s}_i^l) - u_i(y_{i+1}, \boldsymbol{F}_i^l, \boldsymbol{F}_{i+1}^l) + \delta U_{i+1}^l(y_{i+1}),$$

with boundary conditions $U_N^l(y_N) := 0$, $\forall y_N \in \mathcal{Y}'$. An upper bound estimate on the optimal policy value $V_0(y_0, \boldsymbol{F}_0)$ is obtained by averaging the point estimates $U_0^l(y_0)$, $\forall l \in \{1, 2, \ldots, L\}$. We consider the following feasible dual penalties defined using the MLSM value function approximation (17):

$$\hat{V}_{i+1}(y_{i+1}, \boldsymbol{F}_{i+1}; \bar{\beta}_{i+1, y_{i+1}}) - \mathbb{E}\left[\hat{V}_{i+1}(y_{i+1}, \boldsymbol{F}_{i+1}; \bar{\beta}_{i+1, y_{i+1}})|\boldsymbol{F}_i\right]. \tag{35}$$

As alluded to in §5, our choice of basis functions allows the exact computation of the expectation in (35), which is critical to keep the CPU time required for numerically estimating a dual upper bound to a manageable level.

D Additional Numerical Results

Table 4 includes the LB and UB values, their standard errors, and the LB and UB ratios.

Table 4: LB and UB values with standard errors in parentheses and the LB/UB percentages

Year	Month	LB/UB (%)	LB ($/MMBtu)	UB ($/MMBtu)
2011	Jun	99.770	12.596 (0.019)	12.625 (0.001)
2011	Jul	99.731	10.377 (0.019)	10.405 (0.002)
2011	Aug	99.764	11.396 (0.021)	11.423 (0.002)
2011	Sep	99.914	12.778 (0.023)	12.789 (0.002)
2011	Oct	100.015	12.997 (0.023)	12.995 (0.002)
2011	Nov	99.978	13.634 (0.023)	13.637 (0.002)
2011	Dec	99.986	14.758 (0.021)	14.760 (0.002)
2012	Jan	99.949	13.637 (0.019)	13.644 (0.002)
2012	Feb	99.880	14.126 (0.014)	14.143 (0.001)
2012	Mar	99.859	12.715 (0.012)	12.733 (0.001)
2012	Apr	99.917	11.986 (0.011)	11.996 (0.001)
2012	May	99.914	10.463 (0.012)	10.472 (0.001)

References

Ahuja, R. K., T. L. Magnanti, J. B. Orlin. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, New Jersey, USA.

Bertsimas, D., J. N. Tsitsiklis. 1997. *Introduction to Linear Optimization*. Athena Scientific, Belmont, Massachusetts, USA.

Brown, D. B., J. E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4) 1–17.

Heyman, D. P., M. J. Sobel. 2003. *Stochastic Models in Operations Research: Stochastic Optimization*, vol. 2. Dover, Murray Hill, New Jersey.

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