MICHIGAN
ROSS SCHOOL OF BUSINESS

Working Paper

Seasonal Energy Storage Operations with Limited Flexibility

Owen Q. Wu
Stephen M. Ross School of Business
University of Michigan

Derek D. Wang
Stephen M. Ross School of Business
University of Michigan

Zhenwei Qin
DTE Energy
Ann Arbor, Michigan

Ross School of Business Working Paper
Working Paper No. 1160
August 2011

Seasonal Energy Storage Operations with Limited Flexibility

Owen Q. Wu   Derek D. Wang
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan

Zhenwei Qin
DTE Energy, Ann Arbor, Michigan

August 8, 2011

The value of seasonal energy storage depends on how the firm best operates the storage to capture the seasonal price spread. Energy storage operations typically face limited operational flexibility characterized by the speed of storing and releasing energy. A widely used practice-based heuristic, the rolling intrinsic (RI) policy, generally performs well, but can significantly underperform in some cases. In this paper, we aim to understand the gap between the RI policy and the optimal policy, and design improved heuristic policies to close or reduce this gap. A new heuristic policy, the "price-adjusted rolling intrinsic (PARI) policy," is developed based on theoretical analysis of the value of storage options. This heuristic adjusts prices before applying the RI policy, and the adjusted prices inform the RI policy about the values of various storage options. Our numerical experiments show that the PARI policy is especially capable of recovering high value losses of the RI policy. For the instances where the RI policy loses more than 4% of the optimal storage value, the PARI policy on average is able to recover more than 90% of the value loss.

1. Introduction

Energy storage plays an essential role in managing the mismatch between energy supply and demand. Because of the seasonality in demand, energy storage operations exhibit seasonal patterns: Natural gas storage (e.g., depleted reservoir, aquifer) operates in annual cycles; electricity storage (e.g., hydroelectric pumped storage, compressed air storage, batteries) typically has daily cycles.

The value of energy storage depends not only on the seasonal price spread, but also on how the firm best operates the storage. Energy storage operations typically face limited operational flexibility: Firms can choose periods with the best energy prices to buy and sell energy, but the quantities are limited by the storing and releasing capacities, which are determined by physical constraints or contractual terms. Figure 1 shows an example of physical constraints for a typical natural gas storage facility. Panel (a) shows that the maximum injection rate is initially constant and then declines in response to the higher reservoir pressure as working inventory builds up; a reverse trend is observed for the withdrawal rate. (Gas reservoirs hold working gas and base gas. Working inventory refers to gas that can be withdrawn; base gas is needed as permanent inventory to maintain adequate reservoir pressure.) Panel (b) is derived from panel (a) and shows the monthly injection and withdrawal capacities: An empty storage can receive a maximum of 3.1 trillion Btu (British thermal unit) in the first 30 days and less in the following months. It takes about four months to fill up or deplete the storage, or about eight months to complete a cycle.

Figure 1: Injection and withdrawal capacities of a typical natural gas storage facility

Source: Financial Engineering Associates (FEA)

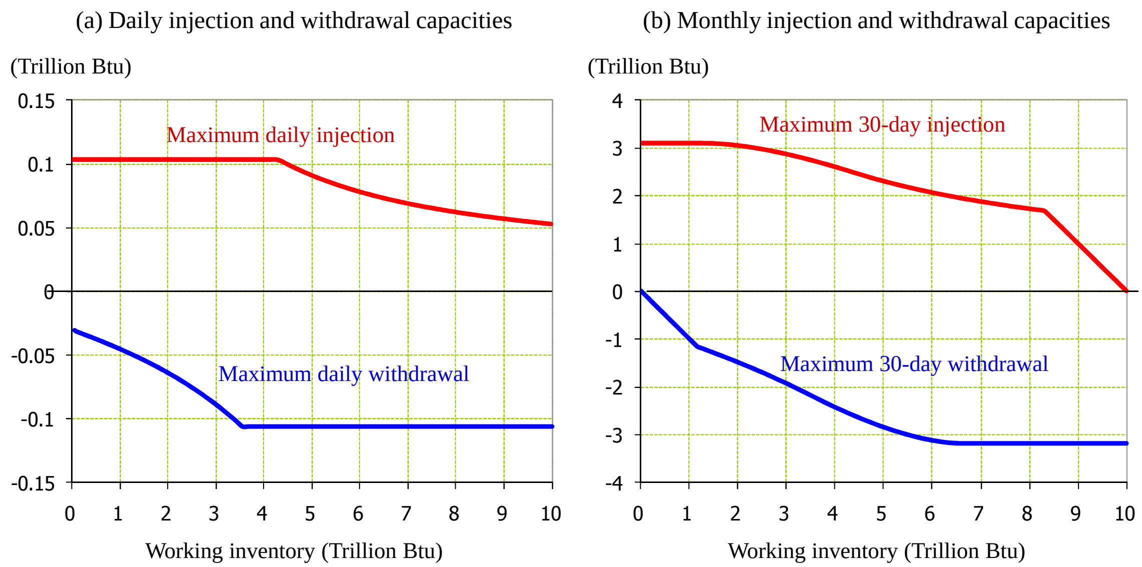

Managing storage with unlimited capacity is known as the warehouse problem, which was first proposed by Cahn (1948). With limited flexibility, storage valuation is considerably more challenging because it involves multiple interacting real options, i.e., options to store or withdraw within capacity limits in every period. Analytical solutions for storage valuation typically do not exist; significant development in numerical techniques of valuation has been seen in recent years, e.g., Manoliu (2004), Chen and Forsyth (2007), Thompson et al. (2009), among others.

In general, finding the optimal storage policy is analytically and computationally challenging. Consequently, heuristic methods have been developed in practice and studied in academia. A widely-used heuristic method is the rolling intrinsic (RI) policy, detailed in Gray and Khandelwal (2004a,b), and is also referred to as the reoptimized intrinsic policy by Secomandi (2010) and Lai et al. (2010). Under the RI heuristic, in each period, the storing or releasing quantity is decided by solving a static optimization problem that involves only forward prices or price forecasts; prices are updated every period and the storage is re-evaluated. The RI policy has near-optimal performance in many circumstances (Secomandi 2010, Lai et al. 2010), but can significantly underperform in some cases.

This paper aims to understand the gap between the RI policy and the optimal policy and to design improved heuristic policies to close or reduce the gap. We design a new heuristic policy called the "price-adjusted rolling intrinsic" (PARI) policy, in which prices are adjusted before applying the RI policy. This simple idea turns out to be very effective: In a three-period problem, the PARI policy is proven to be optimal, and in the multiperiod setting, our numerical results show that the PARI policy is especially capable of recovering high value losses caused by the RI policy.

The price adjustment method is derived based on the understanding of four types of option values in storage operations, briefly described below.

(a) *Value of waiting.* Even if the current price is higher than the expected future prices, it may be beneficial to defer sales when the firm has the flexibility to release energy to capture the expected maximum selling prices.

(b) *Value of avoiding adverse price.* Even if the current price is the lowest compared to the expected future prices, selling some inventory right now may be beneficial because it allows the firm to avoid the expected minimum selling prices in the future.

(c) *Value of counter-seasonal operations.* Price fluctuations may create within-season profit opportunities, which can be captured by counter-seasonal operations, e.g., buying in the selling season.

(d) *Value of raising operational capacity.* When the storing (releasing) speed depends on the inventory level, storing (releasing) less energy in the current period allows the firm to have a higher storing (releasing) capacity in the future to profit from better prices.

The value of waiting and the value of raising operational capacity reduce the firm's incentive to sell, whereas the value of counter-seasonal operations and the value of avoiding adverse price increase that incentive. Thus, it is necessary to strike a balance between these values. We formalize these tradeoffs in this paper.

The rest of this paper is organized as follows. The relevant literature is reviewed in §2. The seasonal storage operations are modeled in §3. The PARI policy is constructed and analyzed in §4 and §5. Numerical results are presented in §6. We conclude the paper with discussion in §7.

2. Literature Review

Managing a fully flexible storage facility is known as the warehouse problem (Cahn 1948). Many researchers have addressed the problem under various settings. The deterministic version of the problem is studied by Charnes and Cooper (1955), Bellman (1956), Prager (1957), and Dreyfus (1957). The warehouse problem with stochastic price variations is considered by Charnes et al. (1966), who find that the optimal policy is a bang-bang type (if the firm acts, it would either fill up the storage or sell all the inventory). Kjaer and Ronn (2008) analyze a model with both spot and futures markets. Hodges (2004) solves a continuous-time model of a fully flexible storage facility.

In practice, storages typically have limited flexibility, due to physical constraints or contractual terms. Secomandi (2010) shows the optimal policy under injection and withdrawal capacities is characterized by two state-dependent basestock targets: If inventory falls between the two targets, it is optimal to do nothing, otherwise the firm should inject or withdraw to bring the inventory as close to the nearer target as possible. In a continuous-time framework, Kaminski, Feng and Pang (2008) prove the optimal policy has a similar structure.

In essence, energy storage operations are multiple interacting real options, that is, options to store or withdraw within capacity limits in every period. Dixit and Pindyck (1994) and Schwartz and Trigeorgis (2001) provide the theoretical background of real options. Analytical valuation of storage options typically do not exist due to the injection and withdrawal constraints. Three computational methods have been developed for storage valuation: numerical partial differential equation techniques (Chen and Forsyth 2007, Thompson, Davison and Rasmussen 2009), binomial/trinormial trees (Manoliu 2004, Parsons 2007), and the Monte Carlo simulation (Boogert and De Jong 2008, Carmona and Ludkovski 2010, Li 2009). Chen and Forsyth (2007) provide a good survey of these computational methods. Our work complements the above works by identifying various types of storage options and revealing useful insights to improve heuristic policies.

Practitioners typically employ two heuristic policies to value seasonal energy storage, the rolling intrinsic (RI) approach and the rolling basket of spread options approach (Gray and Khandelwal 2004a,b, Eydeland and Wolyniec 2003). Lai et al. (2010) refer to them as reoptimized intrinsic value policy and reoptimized linear program policy, respectively. Gray and Khandelwal (2004b, p. 4) state, "Additionally, we have found empirically that, in general, the rolling intrinsic value is equal to the rolling basket value." Lai et al. (2010) employ an approximate dynamic programming approach to value storage with constant capacities and study the effectiveness of the heuristics. They find both heuristics have near-optimal performance. Lai et al. (2011) value the real option to store liquefied natural gas at a regasification terminal. Our work complements the above research by identifying the conditions under which the RI heuristic deviates from the optimal policy and by developing methods to bring the RI heuristic closer to optimality.

3. The Model

Consider an energy storage facility with maximum working inventory level denoted as K. The planning horizon lasts N periods, indexed by $t = 1, 2, \ldots, N$. At the beginning of period t, let x_t be the inventory level in the storage. In this paper, we interchangeably use 'energy level' and 'inventory level,' which are measured in units of energy. The price-taking firm aims to maximize the profit from storage operations.

3.1 Operational Constraints and Costs

Let $\overline{\lambda}(x) \geq 0$ and $\underline{\lambda}(x) \leq 0$ be the capacity functions. Their absolute values, $\overline{\lambda}(x)$ and $-\underline{\lambda}(x)$, express the maximum amount of energy that can be stored and released, respectively, in one period when the period-starting energy level is x. These capacity functions satisfy the following assumption:

Assumption 1 *There exists $H \in (0, K)$ such that $\underline{\lambda}(x) = -x$ when $x \leq H$, and $\underline{\lambda}'(x) \in (-1, 0]$ when $x > H$. There exists $G \in (0, K)$ such that $\overline{\lambda}(x) = K - x$ when $x \geq G$, and $\overline{\lambda}'(x) \in (-1, 0]$ when $x < G$.*

Assumption 1 implies that the storage can be emptied (filled up) within one period if and only if the period-starting inventory level $x \leq H$ $(x \geq G)$. The slopes of the capacity functions imply that the period-ending inventory limits, defined as $\underline{y}(x) \overset{\text{def}}{=} x + \underline{\lambda}(x)$ and $\overline{y}(x) \overset{\text{def}}{=} x + \overline{\lambda}(x)$, are nondecreasing in x.

When the injection and withdrawal speeds are constant for all inventory levels, we have $\underline{\lambda}(x) = \max\{\underline{C}, -x\}$ and $\overline{\lambda}(x) = \min\{\overline{C}, K - x\}$ for some $\underline{C} < 0$ and $\overline{C} > 0$. We refer to this case as the constant capacities case, which is examined by Secomandi (2010) and Lai et al. (2010).

Storing and releasing energy typically involves operational frictions. For example, in natural gas storage operations, the pumps of the storage facility use some of the gas as fuel (Maragos 2002). If q units are to be added to the storage, the firm needs to purchase $(1 + \alpha)q$ units; if q units are withdrawn from the storage, a fraction βq will be lost and $(1 - \beta)q$ can be sold, where α and β are positive constants. In addition to the volume losses, the firm also incurs a variable cost of $c_\alpha q$ when q units are stored, and a variable cost of $c_\beta q$ when q units are withdrawn, where c_α and c_β are positive constants. These costs cover the use of pumps and other equipment (Maragos 2002).

Many firms contract gas storage for one year and must remove the gas before the end of the term (usually March 31, the end of the peak season) or pay a penalty (Buurma 2010). The penalty is typically proportional to the leftover inventory (Carmona and Ludkovski 2010, Chen and Forsyth 2007) or in general form (Boogert and De Jong 2008). We let $p \geq 0$ denote the penalty per unit of inventory at the end of period N; p is realized in period N and may depend on the market prices modeled below.

3.2 Price Model and Problem Formulation

At the beginning of period t, the futures price for delivery in period t is maturing, denoted as $\widetilde{f}_{tt}$. The firm sees this maturing price and other futures prices $\widetilde{f}_{t\tau}$ that mature in period $\tau = t+1, \ldots, N$, and decides the quantity to purchase or sell at price $\widetilde{f}_{tt}$. The settled amount is then stored in or released from the storage over the entire period t.

We make the standard no-arbitrage assumption under which the futures prices are martingales under an equivalent martingale measure Q (see, e.g., Duffie 2001):

$$\widetilde{f}_{t\tau} = \mathsf{E}_t^Q\big[\widetilde{f}_{s\tau}\big], \qquad t < s \leq \tau, \tag{1}$$

where E_t^Q denotes the expectation under Q-measure with information available up to the beginning of period t. If the futures market is absent, all results in this paper continue to hold with $\widetilde{f}_{tt}$ interpreted as the spot price in period t and $\widetilde{f}_{t\tau}$ interpreted as the forecast in period t for the price in period τ. We choose to model the futures market because it provides the firm with instruments to hedge the storage value (perfect hedging is achievable in a complete market).

We refer to $(1+\alpha)\widetilde{f}_{t\tau} + c_\alpha$ as the buying price of inventory, the price the firm must pay for having one unit of inventory available in the storage in period τ. This price includes procurement cost, volume losses, and operating costs. Similarly, we refer to $(1-\beta)\widetilde{f}_{t\tau} - c_\beta$ as the selling price of inventory, which is the net profit the firm obtains from releasing one unit of inventory in period τ.

To derive the expected discounted value of the storage, we note that the expected marked-to-market profit/loss from the futures positions held by the firm is zero under Q-measure, since futures prices are martingales. Hence, if the firm does not have capital constraints, the no-arbitrage value of the storage is the sum of cash flows at maturity dates evaluated under Q-measure and discounted at the risk-free rate (see, e.g., Duffie 2001). Operations of large energy storage facilities often require large sums of capital, thereby increasing the possibility of financial distress during the storing season. Froot and Stein (1998) show that firms require investments to yield a higher return when all risks cannot be frictionlessly hedged. For the purpose of this paper, we assume that the firm discounts the cash flows at a constant rate R. The insights of the paper are intact under any choice of R, including the risk-free rate.

Define $f_{t\tau}$ and $f^b_{t\tau}$ respectively as the selling price and buying price of inventory discounted to the first period:

$$f_{t\tau} \stackrel{\text{def}}{=} e^{-R(\tau-1)}\big[(1-\beta)\widetilde{f}_{t\tau} - c_\beta\big], \qquad f^b_{t\tau} \stackrel{\text{def}}{=} e^{-R(\tau-1)}\big[(1+\alpha)\widetilde{f}_{t\tau} + c_\alpha\big]. \tag{2}$$

Discounting the prices back to the first period allows not to include the discount factor in the problem formulation in (3) below, which simplifies the analytical expressions throughout the paper. Note that for any fixed maturity τ, the discounted selling and buying prices in (2) are still martingales.

Let $\mathbf{f}_t = (f_{t\tau} : \tau = t, t+1, \ldots, N)$ be the *discounted forward selling price curve* (or simply *forward curve* when no confusion arises) observed at the beginning of period t. Let $V_t(x_t, \mathbf{f}_t)$ be the discounted expected profit-to-go from period t onward. Let y_t be the ending inventory in period t, which is decided by the firm at the beginning of period t.

The storage valuation problem can be written as:

$$V_t(x_t, \mathbf{f}_t) = \max_{y_t \in [\underline{y}(x_t), \overline{y}(x_t)]} r(y_t - x_t, f_{tt}) + \mathsf{E}_t^Q \left[V_{t+1}(y_t, \mathbf{f}_{t+1}) \right], \tag{3}$$

where the one-period reward function $r(q, f_{tt}) \stackrel{\text{def}}{=} -f_{tt}^b\, q$, if $q \geq 0$ (purchase), and $r(q, f_{tt}) \stackrel{\text{def}}{=} -f_{tt}\, q$, if $q < 0$ (sell); the period-ending inventory is bounded between $\underline{y}(x_t) = x_t + \underline{\lambda}(x_t)$ and $\overline{y}(x_t) = x_t + \overline{\lambda}(x_t)$. In the last period, the firm sells as much as possible to maximize the profit, and thus,

$$V_N(x_N, f_{NN}) \;=\; -f_{NN}\, \underline{\lambda}(x_N) - \underline{y}(x_N)p. \tag{4}$$

In general, solving the problem in (3)-(4) is complicated. A widely-used heuristic policy is detailed below.

3.3 Rolling Intrinsic Policy

To define the rolling intrinsic (RI) policy, we first define the *intrinsic policy*, a policy that decides in the first period the actions to be performed in each of the remaining periods. The intrinsic policy is found by solving an optimization problem using *only* the forward prices seen in the first period. The corresponding value is called the *intrinsic value*. The RI policy re-optimizes the action in each period by solving the intrinsic valuation problem using the updated forward prices. We refer to the corresponding value as the *rolling intrinsic value*. The RI policy is commonly used in practice (Gray and Khandelwal 2004a,b) and is also referred to as the reoptimized intrinsic policy by Secomandi (2010) and Lai et al. (2010). Because futures prices are martingales, the RI heuristic essentially replaces uncertain prices by their expected values, which is a type of certainty equivalent control studied by Bertsekas (2005). The policy is formally defined below.

Let $V_t^{\mathrm{I}}(x_t, \mathbf{f}_t)$ and $V_t^{\mathrm{RI}}(x_t, \mathbf{f}_t)$ denote the intrinsic value and the rolling intrinsic value of the storage in period t, respectively.

In period t, given the discounted forward selling prices $\mathbf{f}_t = (f_{t\tau} : \tau \geq t)$, the intrinsic value of the storage $V_t^{\mathrm{I}}(x_t, \mathbf{f}_t)$ is determined by:

$$V_N^{\mathrm{I}}(x_N, \mathbf{f}_t) = -f_{tN}\,\underline{\lambda}(x_N) - \underline{y}(x_N)\mathsf{E}_t^Q[p], \tag{5}$$

$$V_s^{\mathrm{I}}(x_s, \mathbf{f}_t) = \max_{y_s \in [\underline{y}(x_s), \overline{y}(x_s)]} r(y_s - x_s, f_{ts}) + V_{s+1}^{\mathrm{I}}(y_s, \mathbf{f}_t), \qquad t \le s < N. \tag{6}$$

When $t = 1$, the recursion in (5)-(6) yields the intrinsic policy in period 1. If the firm implements the intrinsic policy via futures contracts in period 1 and holds all contracts until maturity, then the policy yields the intrinsic value $V_1^{\mathrm{I}}(x_1, \mathbf{f}_1)$.

In the RI policy, the firm solves (5)-(6) in every period with updated forward curve $\mathbf{f}_t$, and adjusts the futures positions accordingly. Let $y_t^\dagger$ be the futures position on the maturing contract in period t, solved from (5)-(6). Then, the rolling intrinsic value of the storage is defined as:

$$V_N^{\mathrm{RI}}(x_N, \mathbf{f}_N) \;=\; V_N^{\mathrm{I}}(x_N, \mathbf{f}_N), \tag{7}$$

$$V_t^{\mathrm{RI}}(x_t, \mathbf{f}_t) \;=\; r(y_t^\dagger - x_t, f_{tt}) + \mathsf{E}_t^Q\left[V_{t+1}^{\mathrm{RI}}(y_t^\dagger, \mathbf{f}_{t+1})\right], \qquad 1 \le t < N. \tag{8}$$

4. Improving the RI Policy: The Three-Period Case

This section introduces the main ideas of improving the RI policy. In §4.1, we consider several simple examples that lead to the construction of a new heuristic policy − the price-adjusted rolling intrinsic (PARI) policy. In §4.2, we prove the optimality of the PARI policy for the three-period setting.

4.1 From RI Policy to PARI Policy

The RI policy solves a deterministic optimization problem every period and may miss potential option values rising from the stochastic evolution of the forward curve. The idea of the PARI policy is to adjust the forward curve to inform the RI policy about the value of various options. The following three examples each illustrate a different option value and introduce a price adjustment scheme to capture the option value.

The common settings of all the examples are as follows. The storage size is $K = 4$ units. The storage can release (store) three units per period as long as inventory (space) is available, i.e., $\underline{\lambda}(x) = \max\{-x, -3\}$ and $\overline{\lambda}(x) = \min\{4 - x, 3\}$. The operating cost parameters are: $\alpha = 2\%$, $\beta = 1\%$, $c_\alpha = c_\beta = \$0.02$. Assume the discount rate $R = 0$. Then, the definitions in (2) imply that $f_{t\tau}^b = \frac{1+\alpha}{1-\beta}(f_{t\tau} + c_\beta) + c_\alpha = 1.03 f_{t\tau} + 0.04$. We assume the storage is initially full and consider a three-period ($N = 3$) selling season problem.

Example 1: Value of waiting. Suppose in period 1 the forward selling price curve is $(f_{11}, f_{12}, f_{13}) =$ (\$5.00, \$4.97, \$4.95). The intrinsic policy can be found by a greedy method: sell three units at the highest price \$5.00 and sell one unit at the second highest price \$4.97. Thus, the intrinsic value of the storage is \$19.97. (Operating costs are accounted for in the selling prices.)

Under the RI policy, the firm first sells three units at \$5.00, as prescribed in the intrinsic policy. In the second period, assume the selling prices (martingales) evolve as follows: $(f_{22}, f_{23}) = (\$5.30,\ \$5.10)$ with probability 0.5, and $(f_{22}, f_{23}) = (\$4.64,\ \$4.80)$ with probability 0.5. Upon price increase, the RI policy is to sell the remaining unit at \$5.30. Upon price decrease, the RI policy is to do nothing in the second period (no incentive to buy because $f^b_{22} = 1.03 \times 4.64 + 0.04 = \$4.82 > f_{23}$) and sell the remaining unit at \$4.80 in the third period. Thus, the remaining unit is sold at an expected price of $(\$5.30 + \$4.80)/2 = \$5.05$. The expected rolling intrinsic value of the storage is \$20.05.

In the above RI policy, the firm effectively sells energy at $\mathsf{E}^Q_1\left[\max\{f_{22}, f_{23}\}\right] = \$5.05$ by exploiting the flexibility of when to sell, but this flexibility is limited: The storage can release at most three units per period. Hence, the optimal policy is to sell one unit at \$5.00 in the first period and sell the remaining three units at \$5.05 in expectation, yielding the optimal expected profit of \$20.15. Thus, although the maturing price f_{11} is the highest on the forward curve, there is a value of delaying sales.

Let us preview one of the key ideas behind the price-adjusted rolling intrinsic (PARI) policy. The original forward curve does not reveal the value of waiting, because $\max\{f_{12}, f_{13}\} < f_{11}$. Suppose we adjust either f_{12} or f_{13} up to \$5.05, and use the adjusted forward curve as the input to the RI policy. Then, because $f_{11} = \$5.00$ is the second highest among the adjusted prices, the RI policy is to sell only one unit at \$5.00. Hence, for this example, adjusting either f_{12} or f_{13} up to $\mathsf{E}^Q_1\left[\max\{f_{22}, f_{23}\}\right]$ informs the RI policy about the value of waiting and brings the RI decision to the optimal. ∎

Example 2: Value of potential purchase. Suppose in period 1 the forward curve is $(f_{11}, f_{12}, f_{13}) = (\$5.00,\ \$4.85,\ \$5.05)$. The intrinsic policy is to sell one unit at \$5.00 and sell the remaining three units at \$5.05, yielding an intrinsic value of \$20.15. Selling more in the first period and buying in the second period cannot improve the intrinsic value, because the buying price $f^b_{12} = 1.03 f_{12} + 0.04 = \$5.04 > f_{11}$.

Under the RI policy, the firm sells one unit in the first period. In the second period, assume the martingale selling prices (f_{22}, f_{23}) is (\$5.20, \$5.20) or (\$4.50, \$4.90) with equal probabilities. If $(f_{22}, f_{23}) = (\$5.20,\ \$5.20)$, the firm sells the remaining three units at \$5.20. If $(f_{22}, f_{23}) = (\$4.50,\ \$4.90)$, the firm faces a low buying price $f^b_{22} = 1.03 f_{22} + 0.04 = \$4.68$ and can make a profit of $f_{23} - f^b_{22} = \$0.22$ per unit by buying at f^b_{22} and selling at f_{23}. However, it can capture this opportunity only if the storage has less than three units at the start of the second period, which is not the case under the RI policy. Hence, the storage value under the RI policy remains \$20.15.

Let us now consider the strategy of selling $1+\varepsilon$ units in the first period, where $\varepsilon \in [0,2]$. Based on Example 1, this strategy gives up some value of waiting: $(\mathsf{E}_1^Q[\max\{f_{22}, f_{23}\}] - f_{11})\varepsilon = \$0.05\varepsilon$, but it brings an extra profit of $\mathsf{E}_1^Q[\max\{f_{23} - f_{22}^b, 0\}]\varepsilon = \$0.11\varepsilon$ from the potential purchase in the second period. The net expected gain is 0.06ε. The optimal policy is to sell three units in the first period, i.e., $\varepsilon = 2$, yielding an extra profit of \$0.12 and raising the storage value to \$20.27.

This leads to the second key idea of the PARI policy. The forward buying price $f_{12}^b = \$5.04$ is too high to reveal the option value of buying inventory in the second period. Let us adjust f_{12}^b down to $\widehat{f}_{12}^b = (\$4.68 + \$5.20)/2 = \$4.94$, implying that f_{12} is lowered to $\widehat{f}_{12} = \$4.76$. Under the adjusted prices $(f_{11}, \widehat{f}_{12}, f_{13}) =$ (\$5.00, \$4.76, \$5.05), the RI policy is to sell three units at the maturing price \$5.00, which coincides with the optimal policy. Note that $f_{13} - \widehat{f}_{12}^b = \$5.05 - \$4.94 = \$0.11$ equals $\mathsf{E}_1^Q[\max\{f_{23} - f_{22}^b, 0\}]$, representing the value of potential purchase. ∎

Example 3: Value of avoiding adverse price. Suppose $(f_{11}, f_{12}, f_{13}) =$ (\$5.00, \$5.05, \$5.02). Note the maturing price f_{11} is the lowest. The intrinsic value is \$20.17, which is the profit of selling three units at \$5.05 and one unit at \$5.02.

The RI policy is to do nothing in the first period. In the second period, assume (f_{22}, f_{23}) is (\$5.40, \$5.10) or (\$4.70, \$4.94) with equal probabilities. Upon price increase (or decrease), the RI policy sells three units at \$5.40 (or \$4.94) and one unit at \$5.10 (or \$4.70). The expected value of the storage under the RI policy is \$20.41.

However, if in the first period the firm sells $\varepsilon \in (0,1]$ units at the lowest price $f_{11} = \$5.00$, then upon price increase (or decrease) it sells $1-\varepsilon$ units at \$5.10 (or \$4.70). Thus, by selling ε units at \$5.00 now, the firm sells ε units less at an expected price $(\$5.10 + \$4.70)/2 = \$4.90$, which equals to the expected minimum price $\mathsf{E}_1^Q[\min\{f_{22}, f_{23}\}]$. The optimal policy is to set $\varepsilon = 1$, and the storage value is improved to \$20.51.

We introduce another idea of the PARI policy that helps the firm avoid selling at the adverse price. The original forward curve does not reveal the adverse price, because $\min\{f_{12}, f_{13}\} > f_{11}$. Suppose we adjust either f_{12} or f_{13} down to $\mathsf{E}_1^Q[\min\{f_{22}, f_{23}\}] = \$4.90$, and use the adjusted forward curve as the input to the RI policy. Then, because $f_{11} = \$5.00$ is no longer the lowest price among the adjusted prices, the RI policy is to sell one unit at \$5.00, which coincides with the optimal policy. Note that $f_{11} - \mathsf{E}_1^Q[\min\{f_{22}, f_{23}\}] = \$0.10$ is exactly the value difference between the optimal policy and the RI policy. ∎

The previous examples show three different option values under constant storing and releasing capacities. In Example 3, if the maximum releasing speed increases in the inventory level, there is an incentive not to sell in the first period, because keeping a higher inventory level raises the releasing capacity in the second period, allowing the firm to sell more at f_{22} and less at f_{23} when $f_{22} > f_{23}$. This is the fourth option value − value of raising operational capacity.

We summarize the four option values in Table 1. For the value of potential purchase, we use a more general term "value of counter-seasonal operations." The third column shows the impact of the option values on the first-period decision. The fourth and fifth columns show the option values and the related spreads seen on the forward curve in the first period.

Table 1: Summary of option values in the selling season

		Impact on y_1^*	Option value	Related spread on forward curve	Price adjustment
$f_{11} > f_{12}$	Value of waiting	$\uparrow$	$\mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big] - f_{11}$	$\max\{f_{12}, f_{13}\} - f_{11}$	$f_{13}\uparrow,\ f_{12}\downarrow$
	Value of counter-seasonal operations	$\downarrow$	$\mathsf{E}_1^Q\big[\max\{f_{23} - f_{22}^b, 0\}\big]$	$f_{13} - f_{12}^b$	
$f_{11} < f_{12}$	Value of avoiding adverse price	$\downarrow$	$f_{11} - \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big]$	$f_{11} - \min\{f_{12}, f_{13}\}$	$f_{13}\downarrow$, f_{12} stays
	Value of raising operational capacity	$\uparrow$	$\mathsf{E}_1^Q\big[\max\{f_{22} - f_{23}, 0\}\big]$	$f_{12} - f_{13}$	

In Table 1, the option values (column 4) typically exceed the corresponding spreads on the forward curve (column 5). The idea of the PARI policy is to adjust the forward curve to bring the deterministic spreads closer to the option values. Interestingly, there exists a set of price adjustments under which the deterministic spreads equal the option values. This set of price adjustments is stated in Definition 1 below; the last column of Table 1 shows the direction of the price adjustments.

Definition 1 Price-adjusted rolling intrinsic (PARI) policy for $N=3$

Step 1. Price adjustment. Based on the forward curve $\mathbf{f}_1$, define a new forward curve $\widehat{\mathbf{f}}_1$ as follows.

(i) When $f_{11} > f_{12}$, define $\widehat{\mathbf{f}}_1 = (f_{11}, \widehat{f}_{12}, \widehat{f}_{13})$ such that

$$\widehat{f}_{12}^b \;=\; \mathsf{E}_1^Q\big[\text{median}\{f_{22}, f_{22}^b, f_{23}\}\big] \qquad \text{and} \qquad \widehat{f}_{13} \;=\; \mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big].$$

(ii) When $f_{11} \le f_{12}$, define $\widehat{\mathbf{f}}_1 = (f_{11}, f_{12}, \widehat{f}_{13})$ where

$$\widehat{f}_{13} \;=\; \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big].$$

Step 2. In the first period, we solve the intrinsic valuation problem (5)-(6) with $\mathbf{f}_1$ replaced by $\widehat{\mathbf{f}}_1$, and implement the corresponding first-period decision.

Step 3. Apply the regular RI policy for the remaining two periods.

The three previous examples assume binomial price processes and constant injection and withdrawal speeds. One surprising result is that the above PARI policy is optimal for the three-period model under general price distributions and capacity functions. We now turn to prove this optimality.

4.2 Optimality of the PARI Policy

We assume the storage can be emptied in two out of three periods, capturing the limited flexibility of typical seasonal storages. Formally, this assumption is stated as follows:

Assumption 2 *(i)* $x_1 > H$. *(ii)* $|\underline{\lambda}(K)| > K - H$.

Part (i) suggests that the initial inventory cannot be sold in a single period. Part (ii) implies that a full storage can release more than $K - H$ in one period. Thus, a full storage can be emptied in two out of three periods. Typical capacity functions satisfying Assumptions 1 and 2 are illustrated in Figure 2. In the figure, H' will be defined in Lemma 1.

Figure 2: Storing and releasing capacity functions for the three-period model

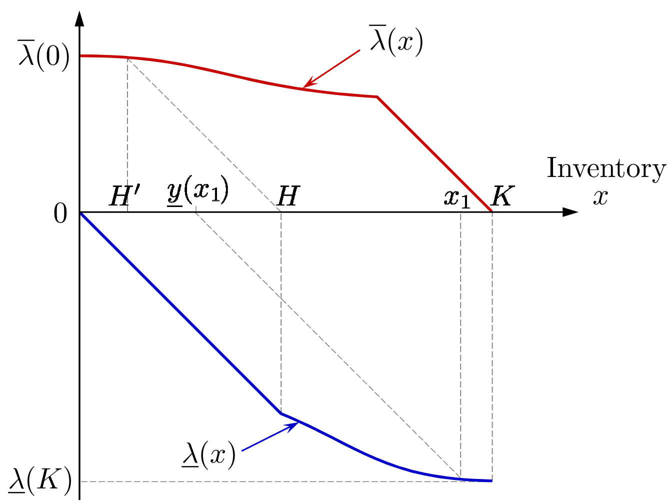

We first show that Step 3 of the PARI policy is optimal for the last two periods.

Proposition 1 *(i) The RI policy is optimal for the last two periods.*
(ii) If the penalty satisfies $\mathsf{P}\Big\{p \geq \dfrac{\underline{s} f_{33} - f_{22}}{1 - \underline{s}}\Big\} = 1$, *where* $\underline{s} \stackrel{def}{=} \sup\{-\underline{\lambda}'(x) : x \in (H, K]\}$, *then for any given first-period decision* y_1, *the second-period RI (optimal) decision is:*

$$y_2^*(y_1, \mathbf{f}_2) - y_1 \ = \begin{cases} \underline{\lambda}(y_1), & \text{if } f_{22} \geq f_{23}, \\ \min\{H - y_1, 0\}, & \text{if } f_{22} < f_{23} \leq f_{22}^b, \\ \min\{H - y_1, \overline{\lambda}(y_1)\}, & \text{if } f_{22}^b < f_{23}. \end{cases} \tag{9}$$

Furthermore, $y_2^*(y_1, \mathbf{f}_2) \leq H$, *and the storage is emptied in the third period.*

The penalty condition in the above proposition is typically satisfied in practice. Under the constant capacities, we have $\underline{s} = 0$ and the penalty condition clearly holds. When the injection and withdrawal speeds vary with inventory, $\underline{s}$ is typically no more than 0.5 (see Figure 1). Thus, the term $\underline{s}f_{33} - f_{22}$ is typically negative, given the fact that the end-of-season selling price f_{33} is typically lower than the mid-season selling price f_{22} (see an example in §6.1).

The RI policy in (9) reacts to the forward curve as follows: If the forward curve is downward sloping $f_{22} \geq f_{23}$, the firm sells as much as possible at price f_{22}. If $f_{22} < f_{23}$, the firm has an incentive to delay sales but needs to sell inventory down to H so that all inventory can be sold in the last period. If the period-starting inventory x_2 is already below H and if the forward curve is steeply upward-sloping $f_{22}^b < f_{23}$, then the firm buys inventory up to or as close as possible to H.

Using the second-period optimal action in (9), we can write the first-period problem as:

$$V_1(x_1, \mathbf{f}_1) = \max_{y_1 \in [\underline{y}(x_1), \overline{y}(x_1)]} U_1(x_1, y_1, \mathbf{f}_1), \tag{10}$$

$$\begin{aligned} U_1(x_1, y_1, \mathbf{f}_1) &= r(y_1 - x_1, f_{11}) + \mathsf{E}_1^Q\big[r\big(y_2^*(y_1, \mathbf{f}_2) - y_1, f_{22}\big) + f_{23}\, y_2^*(y_1, \mathbf{f}_2)\big] \\ &= r(y_1 - x_1, f_{11}) + f_{13}y_1 + \mathsf{E}_1^Q\big[r\big(y_2^*(y_1, \mathbf{f}_2) - y_1, f_{22}\big) + f_{23}(y_2^*(y_1, \mathbf{f}_2) - y_1)\big] \\ &= r(y_1 - x_1, f_{11}) + f_{13}y_1 + \mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[(-f_{22} + f_{23})\underline{\lambda}(y_1) \mid A_1\big] \\ &\quad + \mathsf{P}\{A_2\}\mathsf{E}_1^Q\big[(-f_{22} + f_{23}) \min\{H - y_1, 0\} \mid A_2\big] \\ &\quad + \mathsf{P}\{A_3\}\mathsf{E}_1^Q\big[r\big(\min\{H - y_1, \overline{\lambda}(y_1)\},\, f_{22}\big) + f_{23}(\min\{H - y_1, \overline{\lambda}(y_1)\}) \mid A_3\big], \end{aligned}$$

where $A_1 = \{f_{22} \geq f_{23}\}$ is the downward-sloping forward curve event, $A_2 = \{f_{22} < f_{23} \leq f_{22}^b\}$ is referred to as the slightly upward-sloping forward curve event, and $A_3 = \{f_{22}^b < f_{23}\}$ is the steeply upward-sloping forward curve event.

Next, we prove the optimality of the PARI policy by analyzing the optimal policy and comparing it with the RI policy. We study two cases: $f_{11} > f_{12}$ and $f_{11} < f_{12}$.

4.2.1 Case of $f_{11} > f_{12}$

For this case, we show in the appendix that the problem (10) can be rewritten as:

$$\max_{y_1 \in [\underline{y}(x_1), H]} V^w y_1 + V^c \min\{H - y_1, \overline{\lambda}(y_1)\}, \tag{11}$$

where,

$$V^w \stackrel{\text{def}}{=} \mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big] - f_{11} = \text{ value of waiting}, \tag{12}$$

$$V^c \stackrel{\text{def}}{=} \mathsf{E}_1^Q\big[\max\{f_{23} - f_{22}^b, 0\}\big] \;=\; \text{value of potential purchase (counter-season operations)}. \tag{13}$$

By definition, $V^c \geq 0$, and the sign of V^w is unrestricted. The optimal policy for the first period is summarized in the lemma below. All proofs are included in the online supplement part C.

Lemma 1 *In the first period, if $f_{11} > f_{12}$, then the optimal decision y_1^* is determined as follows:*

(a) If $V^w \leq 0$, then $y_1^* = \underline{y}(x_1)$;

(b) If $V^w > V^c$, then $y_1^* = H$;

(c) If $0 < V^w \leq V^c$, then $y_1^* = \underline{y}(x_1)$ when $\underline{y}(x_1) \geq H'$; when $\underline{y}(x_1) < H'$, y_1^* is determined by $\max_{y_1 \in [\underline{y}(x_1),\, H']} V^w y_1 + V^c \overline{\lambda}(y_1)$, where H' is defined by

$$H' \stackrel{def}{=} \inf\{y \in [0, K] \,:\, y + \overline{\lambda}(y) \geq H\}. \tag{14}$$

The value of waiting V^w and the value of potential purchase V^c drive the decision y_1^* in opposite directions, as shown in Table 1. Lemma 1(b) and (c) reveal the tradeoff between the two values:

- When $V^w > V^c$, the firm should exercise all options of waiting by keeping H units unsold at the end of the first period, leaving no option of purchase in the second period.
- When $0 < V^w < V^c$, the firm should sell as much energy as possible in the first period, as long as it can buy inventory up to H in the second period (this condition is formally stated as $\underline{y}(x_1) \geq H'$, where H' is the level above which the inventory can be raised to H in one period), thereby giving up the options of waiting while maximizing the opportunity of purchase.

Next, we describe the first-period RI policy in the following lemma.

Lemma 2 *In the first period, if $f_{11} > f_{12}$, then under the RI policy, $y_1^\dagger$ is determined as follows:*

(a) If $f_{11} \geq \max\{f_{12}, f_{13}\}$, then $y_1^\dagger = \underline{y}(x_1)$;

(b) If $f_{11} < \min\{f_{12}^b, f_{13}\}$, then $y_1^\dagger = H$;

(c) If $f_{13} > f_{11} \geq f_{12}^b$, then $y_1^\dagger = \underline{y}(x_1)$ when $\underline{y}(x_1) \geq H'$; when $\underline{y}(x_1) < H'$, $y_1^\dagger$ is determined by $\max_{y_1 \in [\underline{y}(x_1),\, H']} (f_{13} - f_{11}) y_1 + (f_{13} - f_{12}^b) \overline{\lambda}(y_1)$.

Comparing the optimal policy and the RI policy, we can prove that if the forward curve in Lemma 2 is adjusted according to Definition 1, the resulting PARI policy is the optimal policy in Lemma 1, as stated in the following proposition.

Proposition 2 *When $N = 3$ and $f_{11} > f_{12}$, the price-adjusted rolling intrinsic (PARI) policy in Definition 1 is optimal. In particular, solving the intrinsic valuation problem (5)-(6) with $\widehat{f}_{12}^b = \mathsf{E}_1^Q[\text{median}\{f_{22}, f_{22}^b, f_{23}\}]$ and $\widehat{f}_{13} = \mathsf{E}_1^Q[\max\{f_{22}, f_{23}\}]$ yields the optimal policy for the first period.*

Raising f_{13} allows the RI policy to see the best selling opportunity in the future, thus capturing the value of waiting. Note that f_{12} is adjusted down because $\widehat{f}^b_{12} = \mathsf{E}^Q_1\big[\text{median}\{f_{22}, f^b_{22}, f_{23}\}\big] \le \mathsf{E}^Q_1\big[\max\{f_{22}, f^b_{22}\}\big] = \mathsf{E}^Q_1\big[f^b_{22}\big] = f^b_{12}$. Lowering f_{12} enlarges the gap between f_{12} and f_{13}, which reflects the value of counter-seasonal operations.

4.2.2 Case of $f_{11} < f_{12}$

The appendix shows that in this case the problem in (10) simplifies to:

$$\max_{y_1 \in [H, \overline{y}(x_1)]} U_1(y_1) = \begin{cases} f_{11}x_1 - V^a y_1 - V^l \underline{\lambda}(y_1), & \text{if } y_1 \in [H, x_1], \\ f^b_{11}x_1 - V^{ab} y_1 - V^l \underline{\lambda}(y_1), & \text{if } y_1 \in (x_1, \overline{y}(x_1)], \end{cases} \tag{15}$$

where

$$V^a \stackrel{\text{def}}{=} f_{11} - \mathsf{E}^Q_1\big[\min\{f_{22}, f_{23}\}\big] = \text{value of avoiding adverse price by selling one more unit}, \tag{16}$$

$$V^{ab} \stackrel{\text{def}}{=} f^b_{11} - \mathsf{E}^Q_1\big[\min\{f_{22}, f_{23}\}\big] = \text{value of avoiding adverse price by buying one less unit}, \tag{17}$$

$$V^l \stackrel{\text{def}}{=} \mathsf{E}^Q_1\big[\max\{f_{22} - f_{23}, 0\}\big] = \text{value of raising operational capacity}. \tag{18}$$

By definition, $V^l \ge 0$, $V^a < V^{ab}$, and the signs of V^a and V^{ab} are unrestricted. Furthermore, $V^a < V^l$ because $V^a - V^l = f_{11} - \mathsf{E}^Q_1\big[\min\{f_{22}, f_{23}\} + \max\{f_{22}, f_{23}\} - f_{23}\big] = f_{11} - f_{12} < 0$.

The following lemma summarizes the optimal policy in this case.

Lemma 3 *In the first period, if $f_{11} < f_{12}$, then the optimal decision y^*_1 is determined as follows:*

(a) If $V^{ab} \le 0$, then $y^*_1 = \overline{y}(x_1)$;

(b) If $V^a \le 0 < V^{ab} \le V^l$, then $y^*_1 \in \underset{y_1 \in [x_1, \overline{y}(x_1)]}{\arg\max} -V^{ab} y_1 - V^l \underline{\lambda}(y_1)$;

(c) If $V^a \le 0 \le V^l < V^{ab}$, then $y^*_1 = x_1$;

(d) If $V^a > 0$, then $y^*_1 \in \underset{y_1 \in [H, \overline{y}(x_1)]}{\arg\max} U_1(y_1)$, where $U_1(y_1)$ is defined in (15)*.*

Example 3 in §4.1 shows that even if $f_{11} < \min\{f_{12}, f_{13}\}$ and all the inventory can be sold in the later periods, selling some inventory in the first period may still be beneficial as it avoids the expected minimum selling price. Similarly, even if $f^b_{11} < \min\{f_{12}, f_{13}\}$, the firm needs to be cautious about buying because the expected minimum price may be below the buying price. We thus refer to V^{ab} in (17) as the value of avoiding adverse price by buying one less unit. Only when $V^{ab} \le 0$, should the firm purchase as much as possible, as confirmed in Lemma 3(a).

The value of avoiding adverse price (V^a or V^{ab}) and the value of raising operational capacity (V^l) drive the decision y^*_1 in opposite directions. When $V^a \le 0$ (implying that selling inventory brings no benefit), the firm trades off between V^l and V^{ab} to decide the purchase quantity, as prescribed in Lemma 3(b) and (c). When $V^a > 0$, the optimal action may be purchase or sell, determined in part (d).

Next, we summarize the first-period RI policy in the following lemma.

Lemma 4 *In the first period, if $f_{11} < f_{12}$, then under the RI policy, $y_1^\dagger$ is determined as follows:*

(a) If $f_{11}^b \le \min\{f_{12}, f_{13}\}$, then $y_1^\dagger = \overline{y}(x_1)$;

(b) If $f_{11} \le \min\{f_{12}, f_{13}\} < f_{11}^b \le f_{12}$, then

$$y_1^\dagger \in \underset{y_1 \in [x_1, \overline{y}(x_1)]}{\arg\max} \; -\big(f_{11}^b - \min\{f_{12}, f_{13}\}\big)\, y_1 - \max\{f_{12} - f_{13}, 0\}\, \underline{\lambda}(y_1);$$

(c) If $f_{11} \le \min\{f_{12}, f_{13}\}$ and $f_{12} < f_{11}^b$, then $y_1^\dagger = x_1$;

(d) If $f_{11} > f_{13}$, then $y_1^\dagger \in \underset{y_1 \in [H, \overline{y}(x_1)]}{\arg\max} \; U_1^{RI}(y_1)$, where

$$U_1^{RI}(y_1) = \begin{cases} f_{11}x_1 - (f_{11} - f_{13})y_1 - (f_{12} - f_{13})\underline{\lambda}(y_1), & \text{if } y_1 \in [H, x_1], \\ f_{11}^b x_1 - (f_{11}^b - f_{13})y_1 - (f_{12} - f_{13})\underline{\lambda}(y_1), & \text{if } y_1 \in (x_1, \overline{y}(x_1)]. \end{cases} \tag{19}$$

We can prove that if the forward curve in Lemma 4 is adjusted according to Definition 1, the resulting PARI policy is the optimal policy in Lemma 3, as stated below.

Proposition 3 *When $N = 3$ and $f_{11} < f_{12}$, the price-adjusted rolling intrinsic (PARI) policy in Definition 1 is optimal. In particular, solving the intrinsic valuation problem* (5)-(6) *with $\widehat{f}_{13} = \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big]$ yields the optimal policy for the first period.*

Adjusting f_{13} alone captures two values. The adjusted price $\widehat{f}_{13}$ informs the firm about the adverse price in the future. Meanwhile, the difference between f_{12} and $\widehat{f}_{13}$ reflects the value of raising operational capacity.

5. Improving the RI Policy: The N-Period Case

In §5.1 and §5.2, we consider a multiperiod model ($N \ge 3$) with constant capacities, and show that the value of waiting, counter-seasonal operations, and avoiding adverse price characterize the optimal policy. Because of the constant capacities, the value of raising operational capacity does not appear in the tradeoffs. In §5.3, we extend the PARI policy to the N-period problem. In §5.4, we further extend the PARI policy to multiple seasons, with each season containing multiple periods.

5.1 Value of Waiting and Value of Avoiding Adverse Price

To focus on the value of waiting and value of avoiding adverse price, we first consider a problem of selling inventory over N periods and delay considering injection (counter-seasonal) operations in §5.2. The capacity functions satisfy the following assumption:

Assumption 3 *(i)* $\underline{\lambda}(x) = \max\{\underline{C}, -x\}$, *where* $\underline{C} < 0$; *(ii)* $K = T|\underline{C}|$ *for some* $T \in \{2, 3, \ldots, N\}$; *(iii)* $\overline{\lambda}(x) = 0$.

Part (i) suggests that the storage can release $|\underline{C}|$ per period until it is empty, following Secomandi (2010) and Lai et al. (2010). Part (ii) assumes that a full storage can be emptied in exactly T periods when releasing energy at the maximum rate. Although part (ii) is not crucial, it simplifies the exposition of our analysis. Part (iii) implies injection operations are not considered.

We let $f_t \equiv f_{tt}$ for notational convenience. For period t, we introduce a T-dimensional vector $\mathbf{u}_t = [u_t^{(1)}, u_t^{(2)}, \ldots, u_t^{(T)}]$, whose k-th element $u_t^{(k)}$ represents the *expected k-th largest price at which inventory may be sold from period t onward*. Formally,

$$\begin{aligned} \mathbf{u}_N &\overset{\text{def}}{=} [\, f_N, 0, \ldots, 0\,], \\ u_t^{(k)} &\overset{\text{def}}{=} k\text{-th largest element of } \left\{f_t,\ \mathsf{E}_t^Q \mathbf{u}_{t+1}\right\}, \qquad k = 1, \ldots, T, \qquad t = 1, \ldots, N-1. \end{aligned} \tag{20}$$

Let $H_k \overset{\text{def}}{=} k|\underline{C}|$, for $k = 0, 1, \ldots, T$. In period $t < N$, when the inventory level is $x_t \in (H_{k-1}, H_k]$, we extend the definitions for the value of waiting and value of avoiding adverse price:

$$V_{tk}^w \overset{\text{def}}{=} \mathsf{E}_t^Q u_{t+1}^{(k-1)} - f_t, \qquad k = 2, \ldots, T, \tag{21}$$

$$V_{tk}^a \overset{\text{def}}{=} f_t - \mathsf{E}_t^Q u_{t+1}^{(k)}, \qquad k = 1, \ldots, T. \tag{22}$$

The optimal policy can be characterized using the values in (21) and (22).

Proposition 4 *Under Assumptions 1 and 3, when $x_t \in (H_{k-1}, H_k]$, $k = 2, \ldots, T$, the optimal decision in period t is as follows:*

$$y_t^* = \begin{cases} \underline{y}(x_t), & \text{if } V_{tk}^w \le 0, \\ H_{k-1}, & \text{if } V_{tk}^w > 0 \text{ and } V_{tk}^a \ge 0, \\ x_t, & \text{if } V_{tk}^a < 0. \end{cases} \tag{23}$$

When $x_t \in (0, H_1]$, $y_t^* = 0$ if $V_{t1}^a \ge 0$, and $y_t^* = x_t$ if $V_{t1}^a < 0$.

Intuitively, when $x_t \in (H_{k-1}, H_k]$, the storage can be emptied in k periods, and the firm aims to sell inventory at the k largest expected prices. When the maturing price f_t is among the $k-1$ highest expected selling prices ($f_t > \mathsf{E}_t^Q u_{t+1}^{(k-1)}$), there is no value of delaying sales ($V_{tk}^w < 0$) and the firm should sell as much as possible, as in the first case of (23).

If the maturing price f_t is lower than the k-th largest expected selling price ($f_t < \mathsf{E}_t^Q u_{t+1}^{(k)}$), then f_t itself is an adverse selling price. Thus, there is no value of avoiding adverse price by selling inventory right now ($V_{tk}^a < 0$), and the firm should do nothing, as in the last case of (23).

When the maturing price f_t is the k-th largest, we have the second case in (23). If the firm sells nothing at f_t, then to sell all inventory it cannot avoid selling some inventory later at a price lower than f_t in expectation. On the other hand, if the firm sells as much as possible right now, then it does not take full advantage of the larger expected selling prices; waiting has a value. The best strategy is to sell down to H_{k-1}, and the remaining H_{k-1} units are expected to be sold at the $k-1$ largest expected selling prices.

The definitions in (21) and (22) are extensions of the definitions of V^w and V^a in (12) and (16), respectively. Note when the storage can be emptied in two out of three remaining periods, i.e., when $N = 3$, $t = 1$, and $k = 2$, (21) and (22) reduce to (12) and (16), respectively.

5.2 Value of Counter-Seasonal Operations

We now allow counter-seasonal operations during the selling season. For ease of illustration, we assume the maximum storing and releasing speeds are the same.

Assumption 4 *(i)* $\overline{\lambda}(x) = \min\{\overline{C}, K - x\}$ *and* $\underline{\lambda}(x) = \max\{\underline{C}, -x\}$*; (ii)* $K = T\overline{C} = T|\underline{C}|$ *for some* $T \in \{2, 3, \ldots, N\}$.

For period t, we introduce a vector $\mathbf{v}_t = [v_t^{(1)}, v_t^{(2)}, \ldots, v_t^{(T)}]$, whose k-th element $v_t^{(k)}$ represents the expected marginal value of inventory in period t when $x_t \in (H_{k-1}, H_k]$. Formally

$$
\begin{aligned}
\mathbf{v}_N &\stackrel{\text{def}}{=} [\, f_N, 0, \ldots, 0\,],\\
v_t^{(k)} &\stackrel{\text{def}}{=} (k+1)\text{-th largest element of } \left\{f_t,\, f_t^b,\, \mathsf{E}_t^Q \mathbf{v}_{t+1}\right\}, \qquad k = 1, \ldots, T, \qquad t = 1, \ldots, N-1. \qquad (24)
\end{aligned}
$$

We inductively prove $\mathbf{u}_t \geq \mathbf{v}_t$. This clearly holds for $t = N$. Suppose $\mathbf{u}_{t+1} \geq \mathbf{v}_{t+1}$. Then, $u_t^{(k)} = k$-th largest element of $\left\{f_t,\, \mathsf{E}_t^Q \mathbf{u}_{t+1}\right\} \geq (k+1)$-th largest element of $\left\{f_t,\, f_t^b,\, \mathsf{E}_t^Q \mathbf{u}_{t+1}\right\} \geq v_t^{(k)}$. We intuitively explain $\mathbf{u}_t \geq \mathbf{v}_t$: Without injection operations, the value of a marginal unit of inventory is the expected price at which this unit can be sold, captured by $\mathbf{u}_t$. When injection is allowed, the marginal unit of inventory brings extra sales revenue but reduces the value of counter-seasonal operations. Hence, $\mathbf{u}_t - \mathbf{v}_t$ indicates the value of counter-seasonal operations.

In period $t \leq N-2$, for $k = 1, \ldots, T$, we define the value of counter-seasonal operations and the value of avoiding adverse price by buying one less unit:

$$
V_{tk}^c \stackrel{\text{def}}{=} \mathsf{E}_t^Q\left[u_{t+1}^{(k)} - v_{t+1}^{(k)}\right], \qquad (25)
$$

$$
V_{tk}^{ab} \stackrel{\text{def}}{=} f_t^b - \mathsf{E}_t^Q u_{t+1}^{(k)}. \qquad (26)
$$

The optimal policy can be characterized by the values defined in (21), (22), (25), and (26).

Proposition 5 *Under Assumptions 1 and 4, when $x_t \in (H_{k-1}, H_k]$, $k = 2, \ldots, T-1$, the optimal decision in period t is as follows:*

$$y_t^* = \begin{cases} \underline{y}(x_t), & \text{if } V_{tk}^w \le V_{t,k-1}^c, \\ H_{k-1}, & \text{if } V_{tk}^w > V_{t,k-1}^c \text{ and } V_{tk}^a + V_{tk}^c \ge 0, \\ x_t, & \text{if } V_{tk}^a + V_{tk}^c < 0 \le V_{tk}^{ab} + V_{tk}^c, \\ H_k, & \text{if } V_{tk}^{ab} + V_{tk}^c < 0 \le V_{t,k+1}^{ab} + V_{t,k+1}^c, \\ \overline{y}(x_t), & \text{if } V_{t,k+1}^{ab} + V_{t,k+1}^c < 0. \end{cases} \quad (27)$$

When $x_t \in (0, H_1]$, the optimal decision is (27) *with the first two cases combined into: $y_t^* = 0$ if $V_{t1}^a + V_{t1}^c \ge 0$. When $x_t \in (H_{T-1}, K]$, the optimal decision is* (27) *with the last two cases combined into: $y_t^* = K$ if $V_{tT}^{ab} + V_{tT}^c < 0$.*

The first three cases in (27) parallel (23). When counter-seasonal operations are not allowed, the optimal policy in (23) considers only the signs of V_{tk}^w and V_{tk}^a. Here in (27), V_{tk}^w and V_{tk}^a are traded off with the value of counter-seasonal operations.

The last two cases in (27) exercise the option of counter-seasonal operations (purchase). The firm should buy as much as possible when buying less provides no combined value of avoiding adverse price and counter-seasonal operations ($V_{t,k+1}^{ab} + V_{t,k+1}^c$). If buying less brings some combined value until inventory hits H_k, then the firm should buy only up to H_k.

The definition of V_{tk}^c in (25) extends that in (13). For the three-period model ($N = 3$), we have:

$$\begin{aligned} u_2^{(1)} - v_2^{(1)} &= \max\{f_{22}, f_{23}\} - \text{median}\{f_{22}, f_{22}^b, f_{23}\} \\ &= \begin{cases} f_{23} - f_{22}^b, & \text{if } f_{22}^b < f_{23} \\ 0, & \text{if } f_{22}^b > f_{23} \end{cases} \\ &= \max\{f_{23} - f_{22}^b,\ 0\}. \end{aligned}$$

Thus, $V_{11}^c = \mathsf{E}_1^Q\left[\max\{f_{23} - f_{22}^b, 0\}\right]$, which is exactly V^c defined in (13).

5.3 *N*-Period PARI Policy

Computing the optimal policy for the multiperiod problem faces the curse of dimensionality, manifested in the recursive definition in (24). In this section, we design a PARI policy for the N-period problem without dramatically increasing the computational burden.

Definition 2 ***N*-period price-adjusted rolling intrinsic (PARI) policy**

Step 1. Set $t = 1$.

Step 2. "Min-Max" price adjustment. Let $f_{t\tau_1}$, $f_{t\tau_2}$, $f_{t\tau_3}$, and $f_{t\tau_4}$ be the maximum, the second *maximum, the second minimum, and the minimum of the futures prices* $\{f_{t\tau} : \tau = t+1, \ldots, N\}$*, respectively. Let* $t' = \tau_1 \wedge \tau_4$*, and* $t'' = \tau_1 \vee \tau_4$*, where* $\wedge$ $(\vee)$ *refers to the* min $($max$)$ *operator.*

(i) When $f_{tt} > f_{tt'}$*, we define* $\widehat{f}_{tt'}$ *and* $\widehat{f}_{tt''}$ *such that*

$$\widehat{f}^{b}_{tt'} = \mathsf{E}^{Q}_{t}\big[\text{median}\{f_{t't'}, f^{b}_{t't'}, f_{t't''}\}\big], \qquad \widehat{f}_{tt''} = \mathsf{E}^{Q}_{t}\big[\max\{f_{\tau_1\wedge\tau_2,\tau_1}, f_{\tau_1\wedge\tau_2,\tau_2}\}\big].$$

(ii) When $f_{tt} \leq f_{tt'}$*, we define* $\widehat{f}_{tt'}$ *and* $\widehat{f}_{tt''}$ *such that*

$$\widehat{f}_{tt'} = f_{tt'}, \qquad \widehat{f}_{tt''} = \mathsf{E}^{Q}_{t}\big[\min\{f_{\tau_3\wedge\tau_4,\tau_3}, f_{\tau_3\wedge\tau_4,\tau_4}\}\big].$$

Step 3. Adjust other prices based on $\widehat{f}_{tt'}$ *and* $\widehat{f}_{tt''}$*. We adjust* $f_{t\tau}$ *by multiplying a scalar that is piecewise linear in* τ*:*

(i) For $t < \tau < t'$*, define* $\widehat{f}_{t\tau} = f_{t\tau}(1 - \delta + \delta \widehat{f}_{tt'}/f_{tt'})$*, where* $\delta = \frac{\tau - t}{t' - t}$*;*

(ii) For $t' < \tau < t''$*, define* $\widehat{f}_{t\tau} = f_{t\tau}\big((1-\delta')\widehat{f}_{tt'}/f_{tt'} + \delta' \widehat{f}_{tt''}/f_{tt''}\big)$*, where* $\delta' = \frac{\tau - t'}{t'' - t'}$*;*

(iii) For $t'' < \tau \leq N$*, define* $\widehat{f}_{t\tau} = f_{t\tau}\big((1-\delta'')\widehat{f}_{tt''}/f_{tt''} + \delta''\big)$*, where* $\delta'' = \frac{\tau - t''}{N - t''}$*.*

Step 4. We solve the intrinsic valuation problem (5)-(6) *with* $\mathbf{f}_t$ *replaced by* $\widehat{\mathbf{f}}_t = (f_{tt}, \widehat{f}_{t,t+1}, \ldots, \widehat{f}_{tN})$*, and implement the decision at the maturing price* f_{tt}*.*

Step 5. If $t < N - 2$*, increase* t *by 1 and go back to Step 2. Otherwise, apply the regular RI policy for the remaining two periods.*

Figure 3: Price adjustment (steps 2 and 3) in the PARI policy

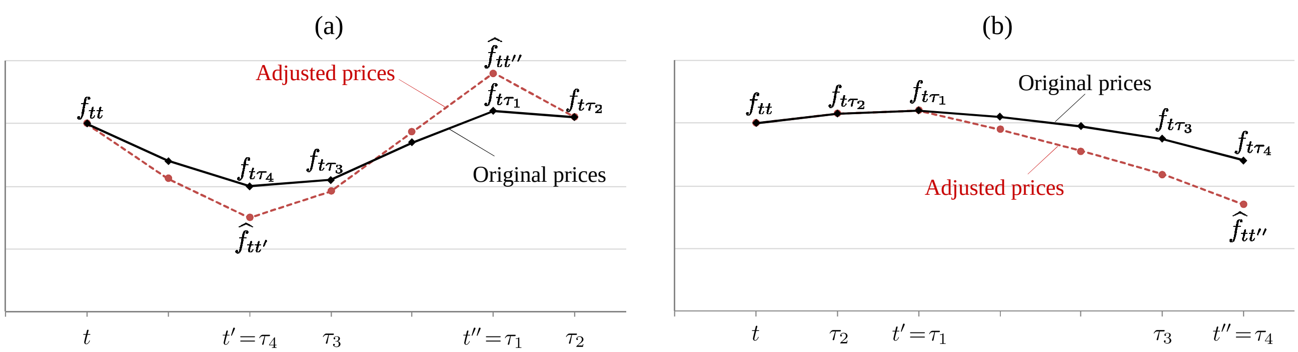

Figure 3 illustrates two typical instances of price adjustment. Step 2 of the above PARI policy resembles the three-period PARI policy. The three focal prices are f_{tt}, $f_{tt'}$, and $f_{tt''}$. The median price formula parallels that in Definition 1, whereas the maximum (minimum) expected selling price is estimated based on the two highest (lowest) futures prices. Note that when $N = 3$, the second maximum price $f_{t\tau_2}$ is the minimum price $f_{t\tau_4}$, and the second minimum price $f_{t\tau_3}$ is the maximum price $f_{t\tau_1}$. Then, the price adjustment formulae in Step 2 are the same as in Definition 1. Indeed, when $N = 3$, the entire policy is identical to that in Definition 1.

The focal prices f_{tt}, $f_{tt'}$, and $f_{tt''}$ divide the forward curve into three segments. Step 3 specifies how each segment should be adjusted if the segment contains prices other than the three focal prices. In essence, the other prices are "attracted" toward $\widehat{f}_{tt'}$ and $\widehat{f}_{tt''}$. This adjustment is important for informing the RI policy about the option values. For example, suppose f_{11} is the highest on the forward curve, the inventory can be sold in two periods, but the optimal policy is not to sell right now. Adjusting $f_{1t''}$ upward in Step 2(i) puts f_{11} in the second highest, which does not stop the RI policy from selling at f_{11}. Step 3 raises other prices, which may signal enough value of waiting such that the RI policy coincides with the optimal policy. Such a heuristic can significantly close the gap between the RI policy and the optimal policy, as will be examined in §6.

Finally, we discuss the computation of the adjusted prices in Step 2. For ease of exposition, assume $\tau_1 < \tau_2$ so that in Step 2(i) we have $\widehat{f}_{tt''} = \mathsf{E}_t^Q\big[\max\{f_{\tau_1\tau_1}, f_{\tau_1\tau_2}\}\big]$. To compute this expectation, we assume $(\log f_{\tau_1\tau_1}, \log f_{\tau_1\tau_2})$ is normally distributed with parameters $(\mu_1, \mu_2, \sigma_1, \sigma_2, \rho)$, where μ_i and σ_i are mean and standard deviation of $\log f_{\tau_1\tau_i}$, $i = 1, 2$, and ρ is the correlation coefficient; these parameters are derived from the forward curve dynamics (see §6.1). Let $f_M = \max\{\log f_{\tau_1\tau_1}, \log f_{\tau_1\tau_2}\}$. Clark (1961) provides the formulae for the moments of the maximum of two normal random variables:

$$\begin{aligned}
\mathsf{E}_t^Q f_M &= \mu_1\Phi(b) + \mu_2\Phi(-b) + a\phi(b),\\
\mathsf{E}_t^Q f_M^2 &= (\mu_1^2 + \sigma_1^2)\Phi(b) + (\mu_2^2 + \sigma_2^2)\mu_2\Phi(-b) + (\mu_1 + \mu_2)a\phi(b),
\end{aligned}$$

where $a^2 = \sigma_1^2 + \sigma_2^2 - 2\sigma_1\sigma_2\rho$, $b = (\mu_1 - \mu_2)/a$, and $\phi(\cdot)$ and $\Phi(\cdot)$ are the probability density function and cumulative distribution function of standard normal random variable, respectively. Clark (1961) also shows that the maximum of two normal random variables is approximately normally distributed. Thus, the adjusted price $\widehat{f}_{tt''}$ can be calculated as

$$\widehat{f}_{tt''} = \mathsf{E}_t^Q \exp(f_M) \approx \exp\big(\mathsf{E}_t^Q f_M + \tfrac{1}{2}\mathsf{Var}_t^Q f_M\big).$$

The expected minimum of two futures prices in Step 2(ii) can be calculated similarly. To estimate $\widehat{f}_{tt'}^b$ in Step 2(i), note that $\text{median}\{f_{t't'}, f_{t't'}^b, f_{t't''}\} = \min\big\{f_{t't'}^b, \max\{f_{t't'}, f_{t't''}\}\big\}$, which can be calculated by repeated use of Clark (1961)'s formulae.

5.4 Multi-Season PARI Policy

Seasonal energy storage operates across seasons. For example, the natural gas industry considers two seasons in storage operation – the withdrawal (peak) season, from November 1 through March 31, and the injection (off-peak) season, from April 1 through October 31 (Energy Information Administration

2011). For storage valuation, we divide the valuation horizon into multiple seasons and apply the PARI policy to each season. Thus, the performance of the PARI policy does not deteriorate when the valuation horizon increases. With distinct price seasonality (e.g., Figure 4 in §6), storage is typically filled during the off-peak season and emptied during the peak season. The off-peak season problem is mathematically equivalent to the peak season problem analyzed in the previous sections, because reducing the inventory level to zero in the peak season is analogous to reducing the space level to zero in the off-peak season. Formally, we define the multi-season PARI policy as follows:

Definition 3 Multi-season PARI policy

Step 1. Divide the planning horizon into a sequence of alternating peak and off-peak seasons. Let N_1 *and* N_2 *be the number of periods in the peak and off-peak seasons, respectively.*

Step 2. Solve peak season problems and off-peak season problems alternately. For each peak season, apply the PARI policy in Definition 2 with $N = N_1$*. For each off-peak season, apply the PARI policy in Definition 2 with* $N = N_2$ *and the following modifications of Step 2:*

(i) When $f^b_{tt} < f^b_{tt'}$*, we define* $\widehat{f}_{tt'}$ *and* $\widehat{f}_{tt''}$ *such that*

$$\widehat{f}_{tt'} = \mathsf{E}^Q_t\big[\text{median}\{f^b_{t't'}, f_{t't'}, f^b_{t't''}\}\big], \qquad \widehat{f}^b_{tt''} = \mathsf{E}^Q_t\big[\min\{f^b_{\tau_3\wedge\tau_4,\tau_3}, f^b_{\tau_3\wedge\tau_4,\tau_4}\}\big].$$

(ii) When $f^b_{tt} \geq f^b_{tt'}$*, we define* $\widehat{f}_{tt'}$ *and* $\widehat{f}_{tt''}$ *such that*

$$\widehat{f}_{tt'} = f_{tt'}, \qquad \widehat{f}^b_{tt''} = \mathsf{E}^Q_t\big[\max\{f^b_{\tau_1\wedge\tau_2,\tau_1}, f^b_{\tau_1\wedge\tau_2,\tau_2}\}\big].$$

In addition, the terminal condition in (5) *is replaced by* $V^I_N(x_N, \mathbf{f}_t) = -f^b_{tN}\,\overline{\lambda}(x_N) + \overline{y}(x_N)p^b$*, where* p^b *is a large constant, which provides incentive to fill up the storage in period* N_2*.*

In the modified (i) above, the buying price $f^b_{tt''}$ is adjusted down to $\widehat{f}^b_{tt''}$ to reflect the value of waiting for a lower buying price, and $f_{tt'}$ is adjusted up to reflect the value of potential sales during the buying season. The price adjustment in (ii) captures the value of avoiding adverse buying price.

6. Application to Natural Gas Storage

6.1 Data and Setup

The average size (for working gas) of a depleted oil/gas reservoir is about 10 trillion Btu (TBtu). We consider a firm leasing a 10 TBtu storage facility for 12 months.

Injection and withdrawal capacities. We consider the case of constant capacities. The capacity pair (injection capacity, withdrawal capacity) takes three values: (2 TBtu/month, 3 TBtu/month), (3 TBtu/month, 4 TBtu/month), and (4 TBtu/month, 5 TBtu/month). Under constant capacities, it is optimal to empty the storage at the end of the horizon regardless of the penalty level (see the proof of Proposition 4). Thus, we set $p = 0$.

Operating cost parameters. For depleted reservoirs, the injection loss rate α is typically between 0% and 3%, the withdrawal loss rate β is between 0% and 2%. Throughout our analysis, we set $\alpha = 1.5\%$, $\beta = 0.5\%$, and the variable operating costs $c_\alpha = c_\beta = \$0.02$ per million Btu. These parameters are consistent with other studies, e.g., Maragos (2002) and Lai et al. (2010).

Discount rate. The discount rate reflects the firm's cost of capital and is typically benchmarked using the London Interbank Offered Rate (LIBOR, available from http://www.liborated.com). We consider three discount rates: 0%, 1%, and 2% above the six-month LIBOR.

Storage contract terms. We consider two different contract terms: (a) the lessee receives an empty storage and returns it empty (such a contract typically starts in April and ends in March); (b) the lessee receives a full storage and returns it full (such a contract typically starts in November and ends in October). These two types of terms are referred to as "seasonal cycling" and "storage carry," respectively, by Eydeland and Wolyniec (2003, p. 354).

Storage valuation under various policies. For the seasonal cycling contracts, the storage value is calculated at the end of March every year for operations from April 1 to March 31. For the storage carry contracts, the value is calculated at the end of October every year. When solving for the optimal policy and the RI policy, we solve the optimization problem without dividing the valuation horizon into peak and off-peak seasons. When implementing the PARI policy, we divide the year into a 7-month off-peak season (April through October) and a 5-month peak season (November through March), and apply the PARI policy in Definition 3.

We value the seasonal cycling contracts in each of the 9 years from 2001-2009, and value the storage carry contracts in each of the 8 years from 2002-2009. At each valuation time, we consider 3 capacity pairs and 3 discount rates. This gives us a total of 153 instances.

Forward curve dynamics. Figure 4 shows the New York Mercantile Exchange (NYMEX) natural gas futures prices observed on the first trading day of March 2005-2009.

We use the NYMEX natural gas futures price data to estimate the following multi-factor martingale model for futures prices (see also Manoliu and Tompaidis (2002) and the references therein):

$$\frac{d\widetilde{f}_{t\tau}}{\widetilde{f}_{t\tau}} = \sum_{j=1}^{n} \sigma_j(t,\tau) dW_j(t), \tag{28}$$

where $W_j(t), j = 1, \ldots, n$, are independent Brownian motions, and $\sigma_j(t,\tau)$ is the volatility of the futures price $\widetilde{f}_{t\tau}$ contributed by the factor j at time t. We employ the principal component analysis

Figure 4: Natural gas forward curve on the first trading day of each March (2005-2009)

Data source: Bloomberg

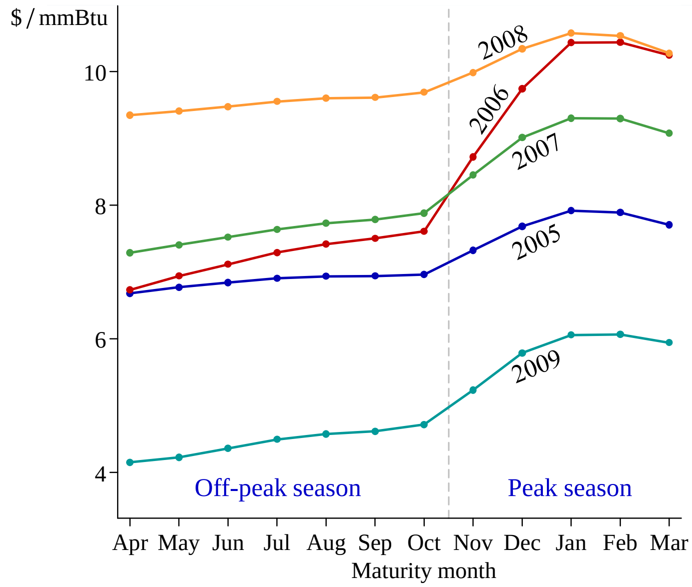

(Basilevsky 1994) to estimate these volatility functions. See Clewlow and Strickland (2000) for examples of principal component analysis for energy prices.

The first two principal components (factors) capture majority of the futures price variations. We build a multi-layer two-factor tree model for the forward curve. Each layer corresponds to a discrete inventory level. This feature is similar to the multi-layer one-factor tree constructed by Jaillet, Ronn and Tompaidis (2004), whereas in our tree each node represents a forward curve. In addition, our tree captures the time-varying volatility feature of the futures prices. The tree construction is described in the online supplement part A.

6.2 Performance of the PARI Policy

We measure the performance of a heuristic policy (RI or PARI policy) by the gap between the storage value under the heuristic policy and optimal storage value, expressed as a percentage of the optimal storage value. Figure 5 compares the percentage storage value losses under the RI policy and PARI policy when valuation is conducted at the end of March (i.e., seasonal cycling contracts). To save space, the results for storage carry contracts are included in the online supplement part B.

The value loss of the PARI policy is remarkably lower than the RI policy. For the 153 instances, the PARI policy achieves an average of 99.8% of the optimal value (minimum 99.13% and maximum 99.99%). That is, the value loss under the PARI is no more than 1% of the optimal value, and 0.2% of the optimal value on average.

Among the 153 cases, there are 5 cases where the RI policy leads to more than 4% value loss in a year, and the PARI policy recovers 92% of that value loss on average. In 13 cases, RI policy results in more than 2% value loss, and the PARI policy recovers 85% of the loss on average. In 26 cases, RI policy loses more than 1% of the value, and the PARI policy recovers 75% of the value loss. For all 153 cases, the PARI policy recovers 64% of the value loss.

Figure 6 depicts this trend over a wider range of percentage value losses. It also shows the

Figure 5: Value loss under RI and PARI policies: Valuation at the end of March

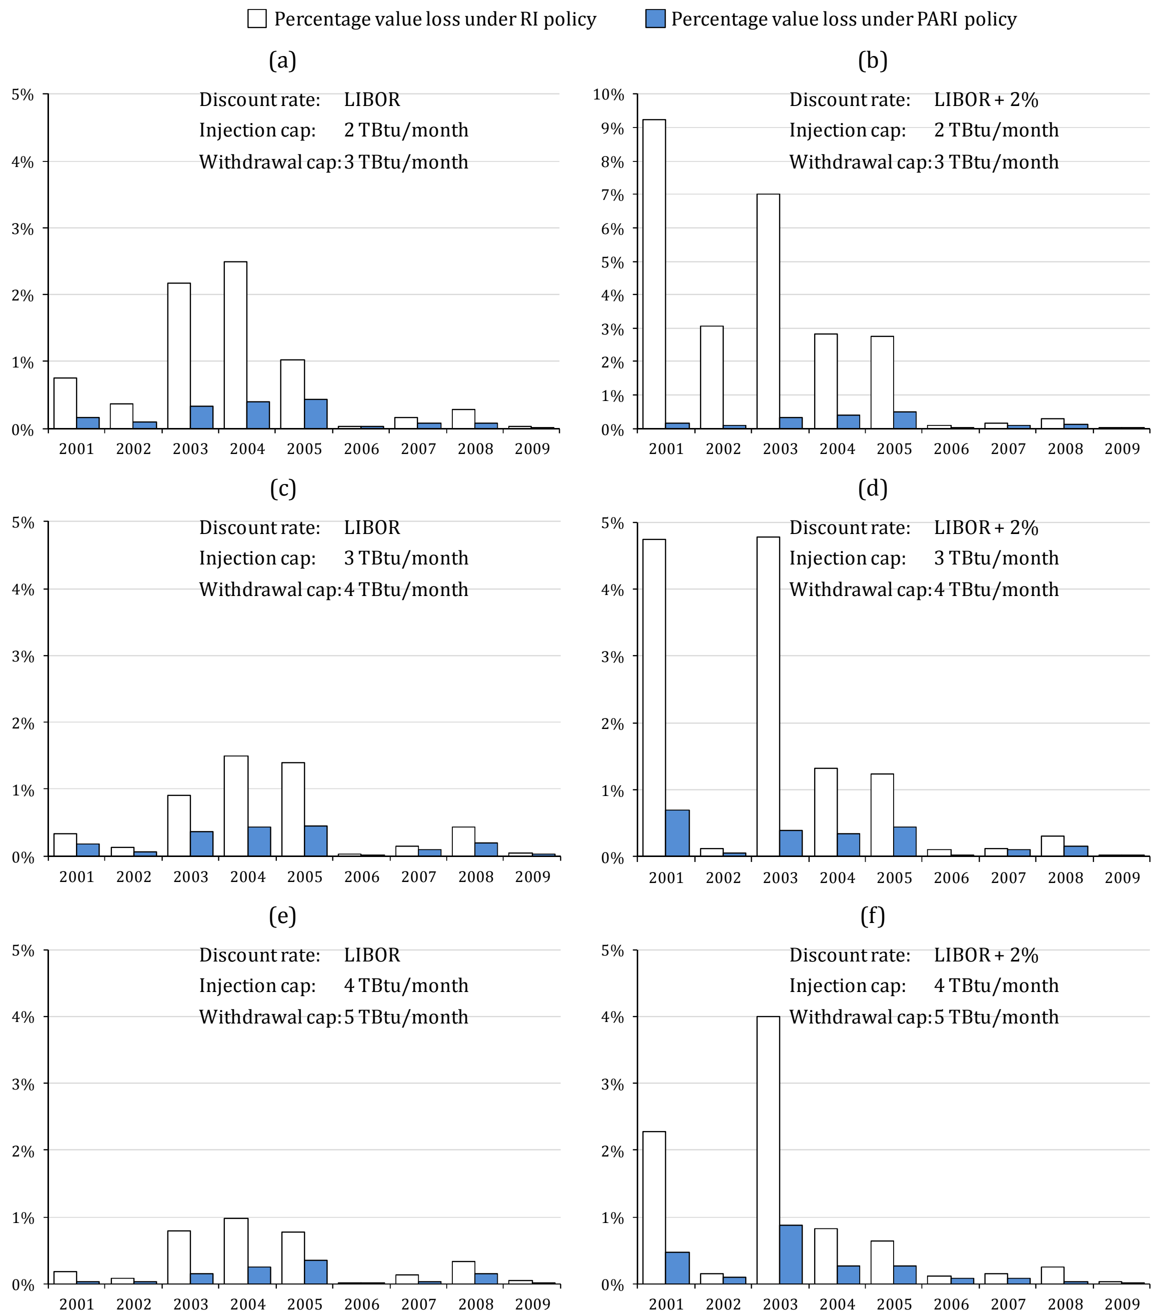

Figure 6: Value loss of the RI policy recovered by the PARI policy

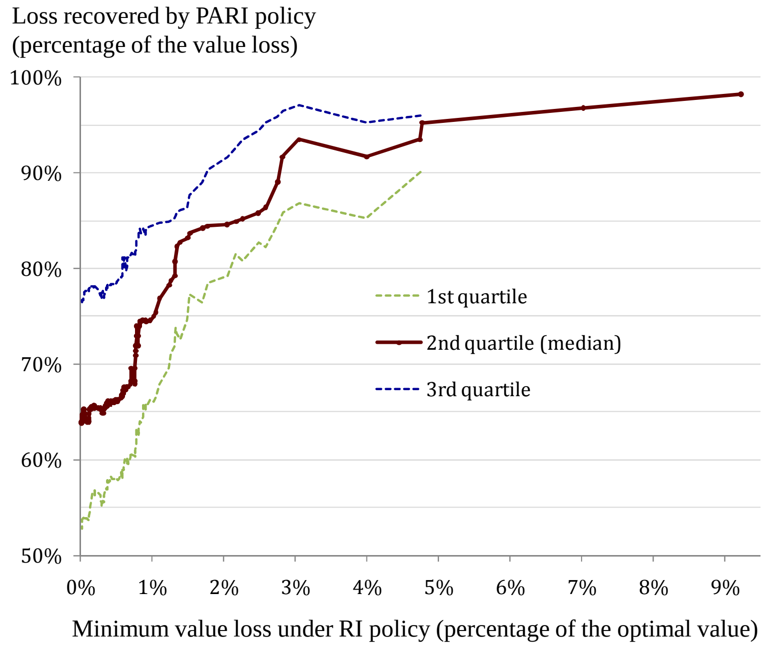

quartiles of the distribution of the loss recovered by the PARI policy (when the RI policy loses more than 5%, there are not enough data points to show the quartiles). Figure 6 suggests the higher the value loss under the RI policy, the more capable the PARI policy in recovering the loss.

We remark on the continuity of the storage value in the discount rate. The discount rate bends the forward curve and affects the option values. The optimal policy takes the option values into

Figure 7: Effect of discount rate on storage value

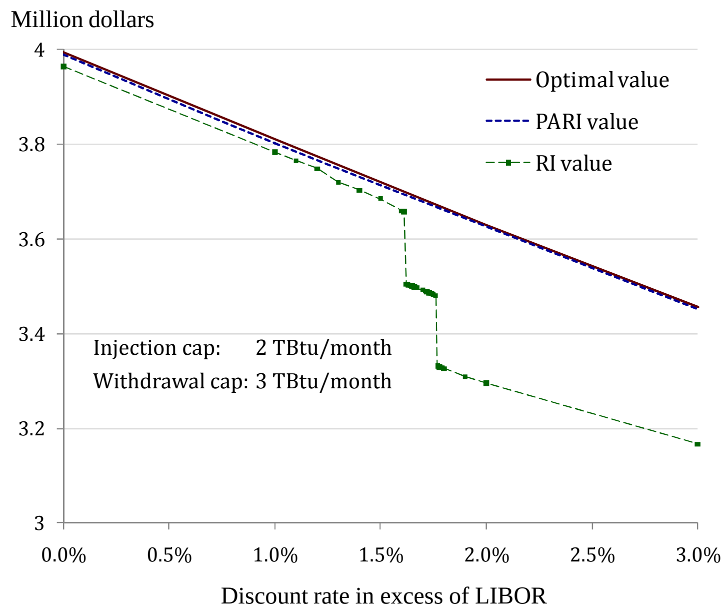

account (e.g., in (11)) and, therefore, the optimal storage value is continuous in the discount rate. However, under the RI policy, a small change in the forward curve can cause the RI policy to miss a lump sum of option values. Thus, the rolling intrinsic value of the storage is, in general, not continuous in the discount rate.

For instance, in Figure 5, for year 2001, the value loss of the RI policy under LIBOR+2% is significantly higher than that under LIBOR. Figure 7 shows how the storage value in 2001 varies with the discount rate. The value loss of the RI policy clearly does not vary smoothly with the discount rate. Remarkably, the PARI policy consistently performs close to the optimal policy. Figure 7 also reinforces the finding in Figure 6 that the PARI policy is especially capable of recovering high value losses of the RI policy.

6.3 Impact of Flexibility

In this section, we study how the operational flexibility of the storage affects the storage value. We vary the flexibility by increasing the injection and withdrawal capacities in tandem, as illustrated in Figure 8(a), which shows capacity functions of the form $\overline{\lambda}(x) = \overline{C} \wedge (10 - x)$ and $\underline{\lambda}(x) = \underline{C} \vee (-x)$, where $|\underline{C}| = \overline{C} + 1$. The storage values are calculated for each capacity function pair indexed by $\overline{C}$.

Figure 8(b) shows when the flexibility increases, the gap between the rolling intrinsic value and

Figure 8: Effect of operational flexibility on storage value

The storage values are calculated in March 2007. Discount rate: LIBOR + 1%

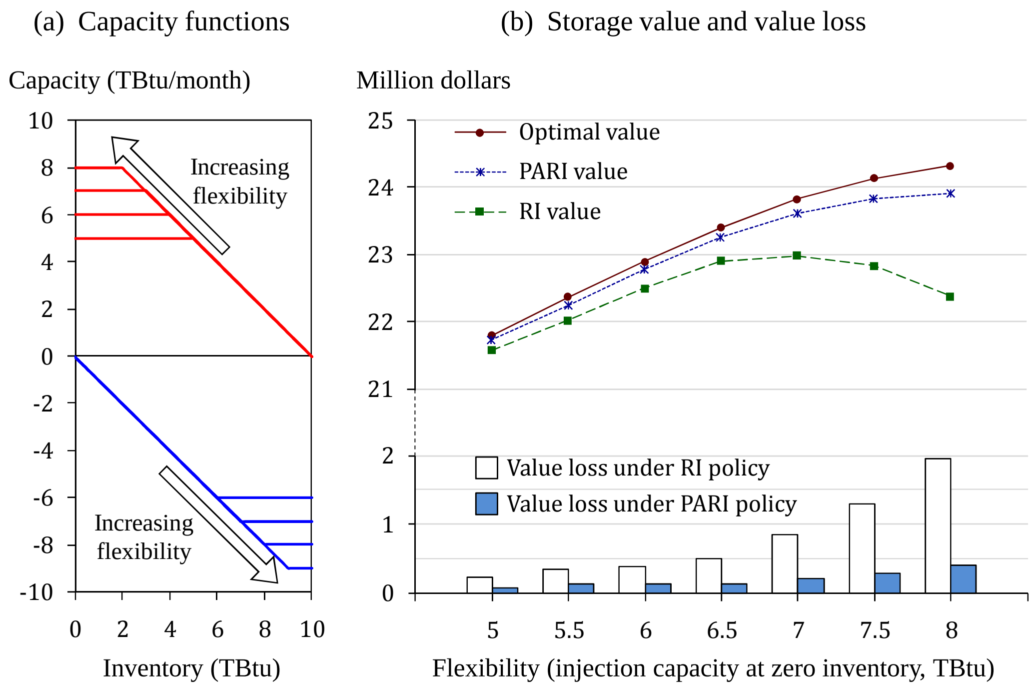

the optimal value widens, and the PARI policy performs significantly better than the RI policy.

One phenomenon is thought-provoking: More flexibility brings more benefits under the optimal policy, but more flexibility may reduce the storage value under the RI policy. In Figure 8(b), the rolling intrinsic value increases and then decreases in flexibility. Intuitively, higher flexibility causes larger deviations of the RI decisions from the optimal decisions, resulting in deteriorating performance. In the online supplement part D, we provide some theoretical support for this finding. We show that if $f_{11} \geq \max\{f_{12}, f_{13}\}$ and $V^w > V^c$, then the expected loss of the RI policy is at least $(V^w - V^c)(H - \underline{y}(x_1))$. If $f_{12} < f_{11} < \min\{f_{12}^b, f_{13}\}$ and $V^w < V^c$, then the expected loss of the RI policy is at least $(V^c - V^w)(H - \max\{\underline{y}(x_1), H'\})$. Note that these lower bounds on the performance gap increase when $\underline{y}(x_1) = x_1 + \underline{\lambda}(x_1)$ decreases or when the releasing capacity $|\underline{\lambda}(x_1)|$ increases. This suggests that more flexibility may cause larger deviation from the optimal policy and lead to higher value loss. Therefore, operational flexibility, if not used with prudence, can be detrimental to the firm. This finding calls for meticulous action to manage relatively flexible storage facilities. The PARI policy does not have the shortcoming of the RI policy: In all of the instances we tested, the storage value under the PARI policy always increases in flexibility.

7. Conclusion and Extensions

Injection and withdrawal capacities are common operational constraints for energy storage facilities. The presence of these constraints renders the optimization of energy storage operations very difficult. In practice, firms use heuristic policies to capture the seasonal price spread under limited flexibility. This paper identifies when and why the rolling intrinsic (RI) policy leads to significant losses and develops an improved heuristic policy called the price-adjusted rolling intrinsic (PARI) policy. The PARI policy is designed based on the analysis of the option values embedded in the optimal policy. Our numerical analysis shows that the gap between the PARI policy and the optimal policy is consistently small, even when the RI policy leads to significant value losses.

Besides natural gas storage, the ideas in this paper and the resulting heuristic policy can be applied to other types of energy storage, such as hydroelectric pumped storage and compressed air energy storage. An interesting future application is the optimization of the battery recharge process for electrical vehicles. Customers may set a time when the battery needs to be fully charged. The electricity distributor aims to meet customers' needs at the minimum procurement cost for energy. This is essentially the problem of filling up the storage with limited flexibility, i.e., the off-peak season problem, with $f_{t\tau}$ interpreted as the price forecast in period t for the price in period τ. We believe the heuristic policies, such as the PARI policy designed in this paper, have great potential to be used in this application.

There are several limitations of this research. First, we do not analyze the combined spot and futures storing and selling strategy. We refer the reader to Goel and Gutierrez (2006), Kjaer and Ronn (2008), and Li (2009) for analysis of models that involve both spot and futures markets. It would be interesting to study how the insights in this paper extend to the setting where both markets are present, and how one can capture the value of spot trading opportunities. Second, we value storage under the forward curve modeled by a two-factor tree. In recent years, the natural gas futures market has seen more variations that cannot be explained by merely two factors. With higher variations, storage options are expected to be more valuable and, therefore, the PARI policy may be more effective in recovering the value loss of the RI policy. Simulation methods can be used in practice to accommodate more factors in the forward curve model. Finally, the firm considered in this paper is a price-taker. The price is determined by the demand and collective behavior of the production and storage firms (see, e.g., Wu and Chen 2010). To consider market equilibrium of storage operations and analyze how energy storage affects energy prices would be another important future direction.

Appendix: Derivation of (11) and (15)

When $f_{11} > f_{12}$, we first show that $y_1^* \le H$. For any policy with $y_1 > H$, we revise that policy by setting $y_1 = H$, while keeping y_2 unchanged (note that $y_2 \le H$ following Proposition 1). The revised policy sells more in the first period and less in the second period. Because $f_{11} > f_{12} = \mathsf{E}_1^Q\big[f_{22}\big]$, the expected profit under the revised policy is higher. Hence, any policy with $y_1 > H$ is sub-optimal, and we must have $y_1^* \le H$. Thus, to solve (10) under $f_{11} > f_{12}$, we need to consider only $y_1 \in [\underline{y}(x_1), H]$. The problem in (10) simplifies to:

$$\max_{y_1 \in [\underline{y}(x_1),\, H]} \quad -f_{11}(y_1 - x_1) + f_{13} y_1 + \mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[f_{22} - f_{23} \mid A_1\big]\, y_1 \\ - \mathsf{P}\{A_3\}\mathsf{E}_1^Q\big[f_{22}^b - f_{23} \mid A_3\big] \min\{H - y_1, \overline{\lambda}(y_1)\}. \tag{29}$$

Ignoring the constant term $f_{11}x_1$, noting that $-\mathsf{P}\{A_3\}\mathsf{E}_1^Q\big[f_{22}^b - f_{23} \mid A_3\big] \equiv \mathsf{E}_1^Q\big[\max\{f_{23} - f_{22}^b, 0\}\big]$, and employing the following identity:

$$f_{13} + \mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[f_{22} - f_{23} \mid A_1\big] = \mathsf{E}_1^Q\big[f_{23}\big] + \mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[f_{22} - f_{23} \mid A_1\big] \\ = \mathsf{E}_1^Q\big[f_{23} + \max\{f_{22} - f_{23}, 0\}\big] \;=\; \mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big], \tag{30}$$

we can rewrite the problem in (29) as:

$$\begin{aligned}\max_{y_1\in[\underline{y}(x_1),H]} \quad & \big(\mathsf{E}_1^Q\big[\max\{f_{22},f_{23}\}\big]-f_{11}\big)y_1+\mathsf{E}_1^Q\big[\max\{f_{23}-f_{22}^b,0\}\big]\min\{H-y_1,\overline{\lambda}(y_1)\}\\ = \;& V^w y_1+V^c\min\{H-y_1,\overline{\lambda}(y_1)\},\end{aligned}$$

which is the problem in (11).

When $f_{11}<f_{12}$, we first show $y_1^*\geq H$. For any policy with $y_1<H$, we can improve the expected profit by raising y_1 to H, i.e., selling $H-y_1\equiv\Delta$ less in the first period and selling Δ more (or buying Δ less) in the second period. The revised policy is feasible, because Assumptions 1 and 2 imply that, by raising y_1 up to H, the releasing capacity $|\underline{\lambda}(x)|$ increases by Δ and the storing capacity $|\overline{\lambda}(x)|$ decreases by at most Δ. Thus, to solve (10) we need to consider only $y_1\geq H$ and the problem in (10) simplifies to:

$$\begin{aligned}\max_{y_1\in[H,\overline{y}(x_1)]} \quad & r(y_1-x_1,f_{11})+f_{13}y_1-\mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[f_{22}-f_{23}\mid A_1\big]\underline{\lambda}(y_1)\\ & +\mathsf{P}\{A_2\cup A_3\}\mathsf{E}_1^Q\big[f_{22}-f_{23}\mid A_2\cup A_3\big](y_1-H).\end{aligned}\tag{31}$$

Using $\mathsf{P}\{A_1\}\mathsf{E}_1^Q\big[f_{22}-f_{23}\mid A_1\big]\equiv\mathsf{E}_1^Q\big[\max\{f_{22}-f_{23},0\}\big]$ and the following identity,

$$f_{13}+\mathsf{P}\{A_2\cup A_3\}\mathsf{E}_1^Q\big[f_{22}-f_{23}\mid A_2\cup A_3\big]=\mathsf{E}_1^Q\big[f_{23}+\min\{f_{22}-f_{23},0\}\big]=\mathsf{E}_1^Q\big[\min\{f_{22},f_{23}\}\big],$$

and ignoring the constant term related to H, we can rewrite the problem in (31) as in (15).

References

Basilevsky, A. 1994. *Statistical Factor Analysis and Related Methods: Theory and Applications*. Wiley. New York.

Bellman, R. 1956. On the theory of dynamic programming – A warehousing problem. *Management Science* **2**(3) 272–275.

Bertsekas, D. P. 2005. *Dynamic Programming and Optimal Control*. Vol. I. 3rd edn. Athena Scientific. Belmont, Massachusetts.

Boogert, A., C. De Jong 2008. Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15**(3) 81–98.

Buurma, C. 2010. Inventories could trump cold for natural-gas prices. *Wall Street Journal*. January 25, 2010.

Cahn, A. S. 1948. The warehouse problem. *Bulletin of the American Mathematical Society* **54** 1073.

Carmona, R., M. Ludkovski 2010. Valuation of energy storage: An optimal switching approach. *Quantitative Finance* **10**(4) 359–374.

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

Charnes, A., W. W. Cooper 1955. Generalizations of the warehousing model. *Operational Research Quarterly* **6**(4) 131–172.

Chen, Z., P. A. Forsyth 2007. A semi-lagrangian approach for natural gas storage valuation and optimal operation. *SIAM Journal on Scientific Computing* **30**(1) 339–368.

Clark, C. E. 1961. The greatest of a finite set of random variables. *Operations Research* **9**(2) 145–162.

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

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

Dreyfus, S. E. 1957. An analytic solution of the warehouse problem. *Management Science* **4**(1) 99–104.

Duffie, D. 2001. *Dynamic Asset Pricing Theory*. 3rd edn. Princeton University Press. Princeton, New Jersey.

Energy Information Administration 2011. Natural gas storage withdrawal season review. Available at http://www.eia.gov/todayinenergy/detail.cfm?id=990.

Eydeland, A., K. Wolyniec 2003. *Energy and Power Risk Management*. John Wiley & Sons. New Jersey.

Financial Engineering Associates 2007. Software module: @ENERGY/Storage. Software information available at http://www.fea.com/products/energy/storage.asp.

Froot, K. A., J. C. Stein 1998. Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics* **47**(1) 55–82.

Goel, A., G. J. Gutierrez 2006. Integrating commodity markets in the optimal procurement policies of a stochastic inventory system. Working paper. University of Texas at Austin.

Gray, J., P. Khandelwal 2004a. Towards a realistic gas storage model. *Commodities Now* (June 2004) 1–4.

Gray, J., P. Khandelwal 2004b. Realistic natural gas storage model II: Trading strategies. *Commodities Now* (September 2004) 1–5.

Hodges, S. D. 2004. The value of a storage facility. Warwick Business School Financial Options Research Centre, Working paper No. 04-142.

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

Kaminski, V., Y. Feng, Z. Pang 2008. Value, trading strategies and financial investment of natural gas storage assets. *Northern Finance Association Conference (Kananaskis, Canada)*. Available at www.northernfinance.org/2008/papers/75.pdf.

Kjaer, M., E. I. Ronn 2008. Valuation of natural gas storage facility. *Journal of Energy Markets* **1**(4) 3–22.

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** 602–616.

Li, Y. 2009. *Natural Gas Storage Valuation*. VDM Verlag. Germany.

Manoliu, M. 2004. Storage options valuations using multilevel trees and calendar spreads. *International Journal of Theoretical and Applied Finance* **7**(4) 425–464.

Manoliu, M., S. Tompaidis 2002. Energy futures prices: Term structure models with Kalman filter estimation. *Applied Mathematical Finance* **9**(1) 21–43.

Maragos, S. 2002. Valuation of the operational flexibility of natural gas storage reservoirs. *In* E. I. Ronn ed. *Real Options and Energy Management*. Risk Books. Chapter 14. pp. 431–456.

Parsons, C. 2007. Valuing commodity storage contracts: A two-factor tree approach. WTM Energy Software LLC. Available at http://ngstoragemodel.com/.

Prager, W. 1957. On warehousing problems. *Operations Research* **5**(4) 504–512.

Schwartz, E. S., L. Trigeorgis 2001. *Real Options and Investment under Uncertainty*. MIT Press. Cambridge, Massachusetts.

Secomandi, N. 2010. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.

Thompson, M., M. Davison, H. Rasmussen 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics* **56**(3) 226–238.

Wu, O. Q., H. Chen 2010. Optimal control and equilibrium behavior of production-inventory systems. *Management Science* **56**(8) 1362–1379.

Online Supplement

A. Two-Factor Tree Model for the Forward Curve Dynamics

This section describes the estimation of forward curve volatility functions from historical data and a two-factor tree model for the price dynamics.

Our historical estimation of forward curve volatility functions follows the principal component analysis (PCA) described in Clewlow and Strickland (2000, §8.6.1). We estimate the volatility functions using the daily futures price data within the three years prior to the date of valuation. For instance, when valuing the storage at the end of March 2005, we use the data from April 2003 to March 2005. The daily futures price data are from Bloomberg.

We construct a two-factor tree model for the evolution of futures prices based on the volatility functions of the first two principal components (factors) that drive the futures price dynamics.

The volatility function for the first factor can be approximated by an exponential function (see, e.g., Clewlow and Strickland 2000): $\sigma_1(t,\tau) = \widehat{\sigma} e^{-\widehat{\theta}(\tau-t)}$, where $\tau - t$ is the time to maturity, and $\widehat{\sigma}$ and $\widehat{\theta}$ are positive constants estimated using a least squares regression: $\ln \sigma_1(t,\tau) = \ln \sigma + \theta(t-\tau) + \varepsilon$.

The exponential volatility function suggests that the volatility increases as a futures contract approaches its maturity. This property of increasing volatility over time can be captured by a tree model with decreasing size of time steps, as shown in Figure A.1. The tree bifurcates at times $t_0(=0), t_1, t_2, \ldots, t_M$. The time step $\Delta t_m \equiv t_{m+1} - t_m$ decreases in m in a certain way described shortly. In each time step prior to the maturity date τ_i of the i-th futures, the price $f_{t\tau_i}$ evolves to either $u_i f_{t\tau_i}$ or $d_i f_{t\tau_i}$. For ease of illustration, Figure A.1 uses only three steps in April. In our actual evaluation, we use many more steps discussed below.

We use the same time steps for all futures contracts, while each futures contract has its own u_i and d_i. Because the first factor drives all futures prices toward the same direction (by different

Figure A.1: Tree model for the first factor

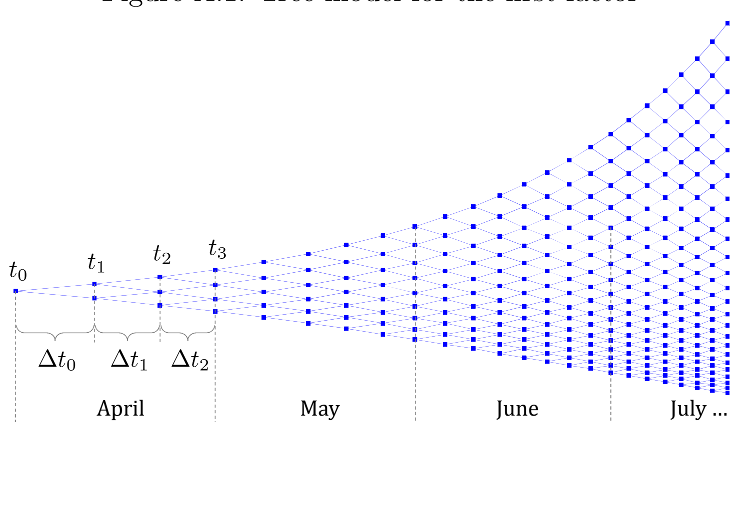

amounts), we must ensure that futures prices move up or down with the same probability. Let the probability of moving up at time t_m be p_m for all futures prices. Matching the first and second moments implied by the binomial tree with those implied by the continuous-time price model, we have:

$$p_m u_i + (1 - p_m) d_i = 1 \tag{A.1}$$

$$p_m u_i^2 + (1 - p_m) d_i^2 = \exp(\sigma_1(t_m, \tau_i)^2 \Delta t_m) = \exp(\widehat{\sigma}^2 e^{2\widehat{\theta}(t_m - \tau_i)} \Delta t_m) \tag{A.2}$$

Note that (A.1) suggests that p_m must be time-invariant because u_i and d_i are constants for each futures contract. This, in turn, suggests that the left side of (A.2) is time-invariant, implying that $\widehat{\sigma}^2 e^{2\widehat{\theta}(t_m - \tau_i)} \Delta t_m$ on the right side must be invariant with respect to m. This specifies how the size of the time steps should shrink over time:

$$\Delta t_{m+1} = e^{-2\widehat{\theta} \Delta t_m} \Delta t_m. \tag{A.3}$$

In our implementation, we set Δt_0 to be 0.4% of a year. Because $\widehat{\theta}$ is estimated at each valuation time, the total number of steps over the 11 months (the last future matures at the beginning of the 12th month) depends on the valuation time. The least number of time steps is 495 (when valuing in March 2004); the maximum number of time steps is 760 (when valuing in March 2002).

We set $p_m = 1/2$ for all m. Then, we can solve for u_i and d_i from (A.2) as follows:

$$u_i = 1 + \sqrt{\exp(\widehat{\sigma}^2 e^{-2\widehat{\theta}\tau_i} \Delta t_0) - 1}, \qquad d_i = 2 - u_i.$$

The volatility function for the second factor $\sigma_2(t, T)$ estimated using PCA generally cannot be approximated by an exponential function, because this factor typically drives the near-term futures and the long-term futures in opposite directions. Consequently, the tree is no longer recombining. To reduce the burden of computing hundreds of instances studied in the paper, we let the tree take one step per month, which leads to $2^{11} = 2048$ nodes at the beginning of the 12th month.

Storage valuation using the above two-factor tree model can be typically solved within 10 minutes with a 2.4GHz Core 2 processor.

B. Storage Carry Contracts

In storage carry contracts, the lessee receives a full storage and returns it full (Eydeland and Wolyniec 2003). Storage carry contracts typically start in November and end in October. We conduct storage valuation at the end of each October for storage operations over the 12 months, starting with a 5-month withdrawal season, followed by a 7-month injection season.

Figure A.2 reports the results, with value in year 2002 referring to the value from November 2001 to October 2002. On average the PARI policy recovers 63% of the value loss. Note that the value loss of the RI policy for the storage carry contracts is lower (less than 2%) compared to the value loss for the seasonal cycling contracts reported in Figure 5. This difference is probably because the peak season forward curve observed at the end of October is typically more curved than the off-peak season forward curve observed at the end of March; the RI policy tends to make suboptimal decisions when the forward curve is flatter.

Figure A.2: Value loss under RI and PARI policies: Valuation at the end of October

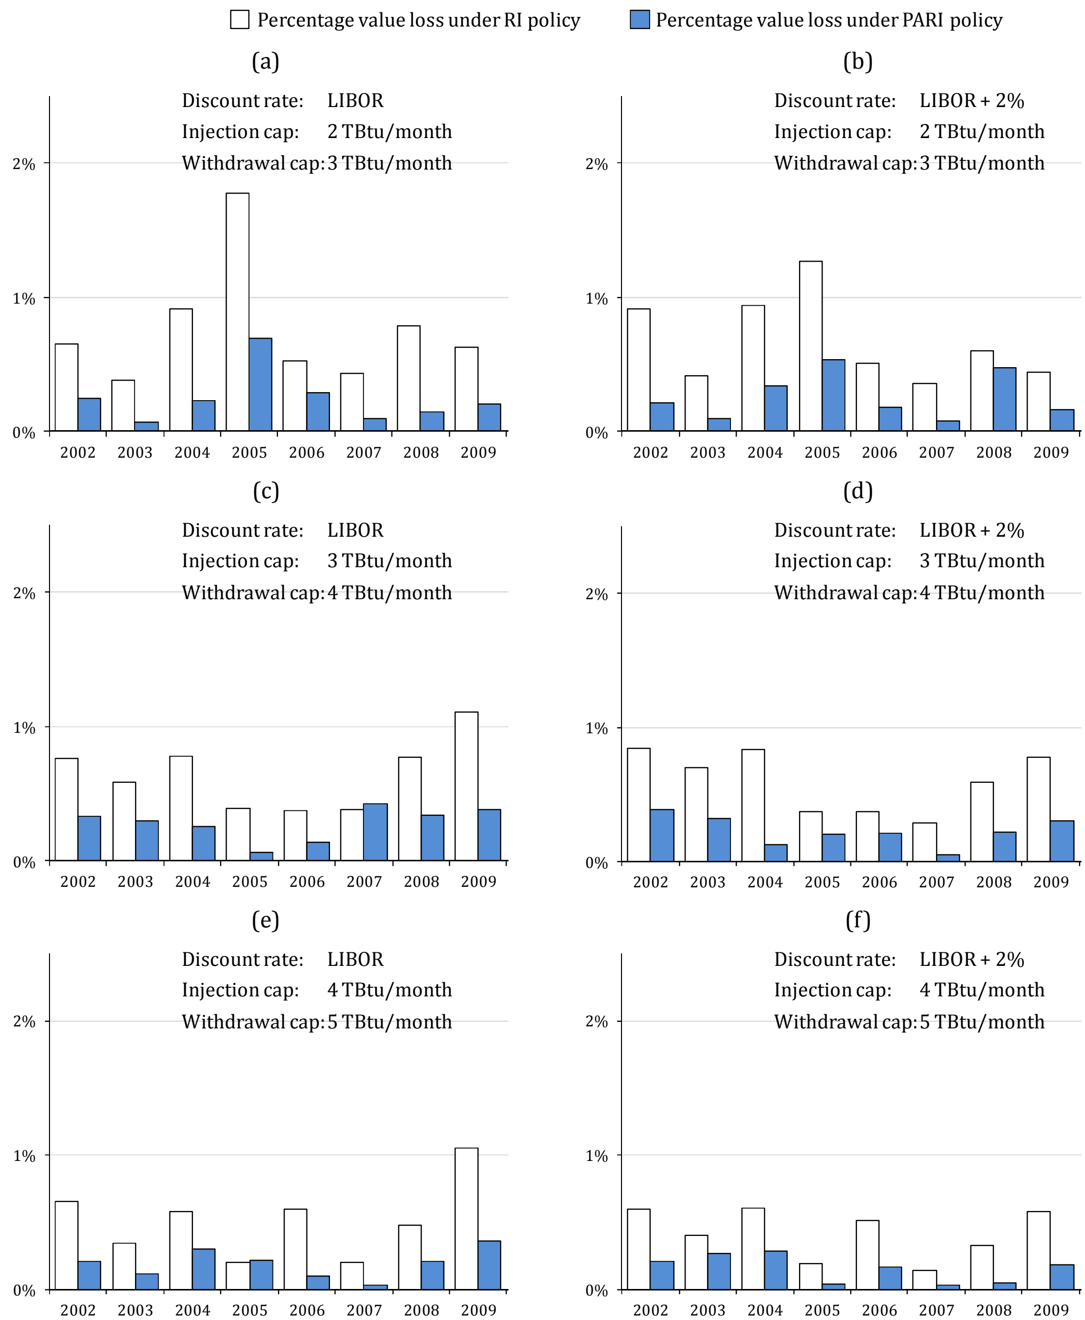

C. Proofs

Proof of Proposition 1. The RI policy is optimal in the last period, because under both policies the firm sells as much as possible to maximize the last-period profit. Next we show that the RI policy is optimal in the second period. Based on (3) and (4), the second-period problem can be written as:

$$V_2(x_2,\mathbf{f}_2) = \max_{y_2\in[\underline{y}(x_2),\overline{y}(x_2)]} U_2(y_2) \stackrel{\text{def}}{=} r(y_2-x_2,f_{22}) + \mathsf{E}_2^Q\big[-f_{33}\underline{\lambda}(y_2)-(y_2+\underline{\lambda}(y_2))p\big], \qquad \text{(A.4)}$$

where, for ease of exposition, we suppress the dependence of $U_2(y_2)$ on x_2 and $\mathbf{f}_2$.

Based on (5) and (6), the second-period RI policy is determined by:

$$\max_{y_2\in[\underline{y}(x_2),\overline{y}(x_2)]} r(y_2-x_2,f_{22}) - f_{23}\underline{\lambda}(y_2)-(y_2+\underline{\lambda}(y_2))\mathsf{E}_2^Q[p],$$

which is identical to (A.4), noting the martingale property of f_{t3}. Thus, the RI policy is optimal in the second period. Next, we prove that the optimal policy has the form in (9).

For $y_2\in(H,K]$ and $y_2\neq x_2$, the first-order derivative of the objective in (A.4) is

$$\begin{aligned} U_2'(y_2) &= \partial r(y_2-x_2,f_{22})/\partial y_2 - \underline{\lambda}'(y_2)f_{23} - (1+\underline{\lambda}'(y_2))\mathsf{E}_2^Q[p] \\ &\le -f_{22} + \underline{s}f_{23} - (1-\underline{s})\mathsf{E}_2^Q[p] \;\le\; 0. \end{aligned}$$

The first inequality follows from two facts: The definition of $r(q,f_{22})$ implies $\partial r(q,f_{22})/\partial q \le -f_{22}$, and the definition of $\underline{s}$ leads to $-\underline{\lambda}'(y_2)\le \underline{s}$, for $y_2\in(H,K]$. The last inequality is because the condition $\mathsf{P}\{p\ge \frac{\underline{s}f_{33}-f_{22}}{1-\underline{s}}\}=1$ implies $\mathsf{E}_2^Q[p]\ge \frac{\underline{s}f_{23}-f_{22}}{1-\underline{s}}$.

Because $U_2'(y_2)\le 0$ for $y_2\in(H,K]$, we need to consider only $y_2\le H$ in solving (A.4). Assumption 1 implies $\underline{\lambda}(y_2)=-y_2$ when $y_2\le H$. Thus, the problem in (A.4) becomes

$$V_2(x_2,\mathbf{f}_2) = \max_{y_2}\big\{r(y_2-x_2,f_{22})+f_{23}y_2 \;:\; \underline{y}(x_2)\le y_2\le \min\{H,\overline{y}(x_2)\}\big\}.$$

The solution to the above problem is:

$$y_2^*(x_2,\mathbf{f}_2) \;=\; \begin{cases} \underline{y}(x_2), & \text{if } f_{22}\ge f_{23},\\ \min\{H,\,x_2\}, & \text{if } f_{22}<f_{23}\le f_{22}^b,\\ \min\{H,\,\overline{y}(x_2)\}, & \text{if } f_{22}^b<f_{23}, \end{cases}$$

which leads to the optimal decision expressed in (9) in the paper. ∎

Proof of Lemma 1. Consider the objective (11) in the paper:

$$\max_{y_1\in[\underline{y}(x_1),H]} V^w y_1 + V^c\min\{H-y_1,\overline{\lambda}(y_1)\}.$$

Under Assumption 1, $\min\{H-y_1,\overline{\lambda}(y_1)\}$ is decreasing in y_1 at a rate no faster than the unit rate. (a) Since $V^c\ge 0$ by definition, the second term in the objective (11) is decreasing in y_1. When

$V^w \le 0$, the first term is also decreasing in y_1 and, therefore, the optimal solution is $y_1^* = \underline{y}(x_1)$.

(b) When $V^w > V^c$, the objective can be written as:

$$(V^w - V^c)y_1 + V^c\big(y_1 + \min\{H - y_1, \overline{\lambda}(y_1)\}\big) \;=\; (V^w - V^c)y_1 + V^c \min\{H, \overline{y}(y_1)\},$$

which is increasing in y_1, because $\overline{y}(y_1)$ is nondecreasing in y_1. Hence, $y_1^* = H$.

(c) When $0 < V^w \le V^c$, the objective can be written as:

$$(V^w - V^c)y_1 + V^c \min\{H, \overline{y}(y_1)\} \;=\; \begin{cases} (V^w - V^c)y_1 + V^c H, & \text{if } y_1 \ge H', \\ V^w y_1 + V^c\overline{\lambda}(y_1), & \text{if } y_1 < H'. \end{cases}$$

Thus, the objective is decreasing in y_1 for $y_1 \ge H'$. If $\underline{y}(x_1) \ge H'$, then $y_1^* = \underline{y}(x_1)$. If $\underline{y}(x_1) < H'$, then $y_1^* \in [\underline{y}(x_1),\, H']$ and is determined by maximizing $V^w y_1 + V^c\overline{\lambda}(y_1)$. ∎

Proof of Lemma 2. Had we not known the optimal policy, we would prove Lemma 2 from scratch. With the optimal policy derived in Lemma 1, a short-cut is available. If we set the volatilities of futures prices to be zero, then the optimal policy in Lemma 1 becomes the RI policy. Specifically, under the zero price volatilities assumption, (12)-(13) in the paper become

$$V^w = \max\{f_{12}, f_{13}\} - f_{11} \qquad \text{and} \qquad V^c = \max\{f_{13} - f_{12}^b, 0\}.$$

We now show that each part of Lemma 1 becomes the corresponding part of Lemma 2:

(a) $V^w \le 0$ is equivalent to $f_{11} \ge \max\{f_{12}, f_{13}\}$.

(b) Because $f_{11} > f_{12}$, $V^w > V^c$ is equivalent to $f_{13} - f_{11} > \max\{f_{13} - f_{12}^b, 0\}$ or $f_{11} < \min\{f_{12}^b, f_{13}\}$.

(c) Based on the equivalence in (a) and (b) above, we can see that $0 < V^w \le V^c$ is equivalent to $f_{13} > f_{11} \ge f_{12}^b$. The maximization problem in Lemma 1(c) is also equivalent to that in Lemma 2(c) because $V^w = f_{13} - f_{11}$ and $V^c = f_{13} - f_{12}^b$. ∎

Proof of Proposition 2. The price is adjusted such that $\widehat{f}_{12}^b = \mathsf{E}_1^Q\big[\text{median}\{f_{22}, f_{22}^b, f_{23}\}\big]$ and $\widehat{f}_{13} = \mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big]$. Using $\widehat{\mathbf{f}}_1 = (f_{11}, \widehat{f}_{12}, \widehat{f}_{13})$ as the input prices of the RI policy, we show that each part of Lemma 2 is equivalent to the corresponding part in Lemma 1:

(a) $f_{11} \ge \max\{\widehat{f}_{12}, \widehat{f}_{13}\} = \widehat{f}_{13} = \mathsf{E}_1^Q\big[\max\{f_{22}, f_{23}\}\big]$ is equivalent to $V^w \le 0$.

(b) $f_{11} < \min\{\widehat{f}_{12}^b, \widehat{f}_{13}\} = \widehat{f}_{12}^b = \mathsf{E}_1^Q\big[\text{median}\{f_{22}, f_{22}^b, f_{23}\}\big]$ is equivalent to $V^w > V^c$, because

$$\text{median}\{f_{22}, f_{22}^b, f_{23}\} = \max\{f_{22}, f_{23}\} - \max\{f_{23} - f_{22}^b, 0\}. \tag{A.5}$$

One can verify (A.5) by considering three cases: $f_{22} < f_{22}^b < f_{23}$, $f_{22} < f_{23} < f_{22}^b$, and $f_{23} < f_{22} < f_{22}^b$.

(c) Based on the equivalent relations in (a) and (b), $\widehat{f}_{13} > f_{11} \ge \widehat{f}_{12}^b$ is equivalent to $0 < V^w \le V^c$. Furthermore, the maximization problem in Lemma 2(c) is identical to that in Lemma 1(c) because $V^w = \widehat{f}_{13} - f_{11}$ and $V^c = \widehat{f}_{13} - \widehat{f}_{12}^b$, where the latter is due to (A.5). ∎

Proof of Lemma 3. Consider the objective (15) in the paper:

$$\max_{y_1\in[H,\overline{y}(x_1)]} U_1(y_1) = \begin{cases} f_{11}x_1 - V^a y_1 - V^l \underline{\lambda}(y_1), & \text{if } y_1 \in [H, x_1], \\ f_{11}^b x_1 - V^{ab} y_1 - V^l \underline{\lambda}(y_1), & \text{if } y_1 \in (x_1, \overline{y}(x_1)]. \end{cases}$$

(a) Because $V^l \geq 0$ by definition and $\underline{\lambda}(y_1)$ is decreasing in y_1 under Assumption 1, the term $-V^l\underline{\lambda}(y_1)$ in the objective is increasing in y_1. When $V^{ab} \leq 0$, the terms $-V^a y_1$ and $-V^{ab}y_1$ are also increasing in y_1 and, therefore, the optimal solution is $y_1^* = \overline{y}(x_1)$.

(b) When $V^a \leq 0 < V^{ab}$, $U_1(y_1)$ is increasing for $y_1 \in [H, x_1]$, and the optimal decision is determined by maximizing $-V^{ab}y_1 - V^l\underline{\lambda}(y_1)$ for $y_1 \in [x_1, \overline{y}(x_1)]$.

(c) Continue from part (b). If $V^l < V^{ab}$, then the maximizer of $-V^{ab}y_1 - V^l\underline{\lambda}(y_1)$ is $y_1^* = x_1$.

(d) When $V^a > 0$, the objective is not monotone in general and the optimal solution may lie anywhere between H and $\overline{y}(x_1)$. ∎

Proof of Lemma 4. Parallel to the proof of Lemma 2, when the price volatilities are assumed to be zero, the optimal policy in Lemma 3 becomes the RI policy stated in this lemma. ∎

Proof of Proposition 3. The adjusted price is $\widehat{\mathbf{f}}_1 = (f_{11}, f_{12}, \widehat{f}_{13})$, where $\widehat{f}_{13} = \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big]$. Note the following relations:

$$f_{11} - \widehat{f}_{13} = f_{11} - \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big] = V^a, \tag{A.6}$$

$$f_{11}^b - \widehat{f}_{13} = f_{11}^b - \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big] = V^{ab}, \tag{A.7}$$

$$f_{12} - \widehat{f}_{13} = f_{12} - \mathsf{E}_1^Q\big[\min\{f_{22}, f_{23}\}\big] = \mathsf{E}_1^Q\big[\max\{f_{22} - f_{23}, 0\}\big] = V^l, \tag{A.8}$$

$$f_{11}^b - f_{12} = (f_{11}^b - \widehat{f}_{13}) - (f_{12} - \widehat{f}_{13}) = V^{ab} - V^l. \tag{A.9}$$

Using $\widehat{\mathbf{f}}_1 = (f_{11}, f_{12}, \widehat{f}_{13})$ as the input prices of the RI policy, we show that each part of Lemma 4 is equivalent to the corresponding part in Lemma 3:

(a) $f_{11}^b \leq \min\{f_{12}, \widehat{f}_{13}\} = \widehat{f}_{13}$ is equivalent to $V^{ab} \leq 0$, due to (A.7).

(b) $f_{11} \leq \min\{f_{12}, \widehat{f}_{13}\} = \widehat{f}_{13} < f_{11}^b \leq f_{12}$ is equivalent to $V^a \leq 0 < V^{ab} \leq V^l$ due to (A.6), (A.7), and (A.9). The maximization problem in Lemma 4(b) is identical to that in Lemma 3(b) because $f_{11}^b - \min\{f_{12}, \widehat{f}_{13}\} = f_{11}^b - \widehat{f}_{13} = V^{ab}$ and $\max\{f_{12} - \widehat{f}_{13}, 0\} = f_{12} - \widehat{f}_{13} = V^l$.

(c) $f_{11} \leq \min\{f_{12}, \widehat{f}_{13}\} \leq f_{12} < f_{11}^b$ is equivalent to $V^a \leq 0 \leq V^l < V^{ab}$ due to (A.6) and (A.9).

(d) $f_{11} > \widehat{f}_{13}$ is equivalent to $V^a > 0$ due to (A.6). The maximization problem in Lemma 4(d) is identical to that in Lemma 3(d), because (A.6)-(A.8) imply that the objective in (19) is identical to the objective in (15). ∎

Proof of Proposition 4. The multiperiod problem is formulated in (3)-(4), and simplified below under Assumption 3.

$$V_t(x_t, \mathbf{f}_t) = \max_{y_t \in [\underline{y}(x_t),\, x_t]} (x_t - y_t) f_t + \mathsf{E}_t^Q\big[V_{t+1}(y_t, \mathbf{f}_{t+1})\big], \tag{A.10}$$

$$V_N(x_N, \mathbf{f}_N) = -f_N\, \underline{\lambda}(x_N). \tag{A.11}$$

Note that under constant capacities, there is no value of raising withdrawal capacity by withholding sales. Thus, it is optimal to empty the storage by the end of period N, and the penalty term is not needed in (A.11). Formally, we show that $y_t^* \leq (N-t)|\underline{C}|$ and, in particular, $y_N^* = 0$. If $y_t > (N-t)|\underline{C}|$, then for the remaining $N-t$ periods, the best policy is to sell $|\underline{C}|$ every period, leaving $\big(y_t - (N-t)|\underline{C}|\big)$ units unsold in the last period. Thus, $y_t > (N-t)|\underline{C}|$ is a suboptimal decision.

We now inductively prove that for any $\mathbf{f}_t$, $V_t(x_t, \mathbf{f}_t)$ is a concave piece-wise linear function in x_t with slope $u_t^{(k)}$ defined in (20) for $x_t \in (H_{k-1}, H_k]$, $k = 1, \dots, T$.

First, because $\underline{\lambda}(x)$ has slope -1 for $x \in (0, H_1]$ and zero slope otherwise, $V_N(x_N, \mathbf{f}_N)$ is concave in x_N and has slope $u_N^{(k)}$ for $x_t \in (H_{k-1}, H_k]$. Suppose $V_{t+1}(y_t, \mathbf{f}_{t+1})$ is concave in y_t with slope $u_{t+1}^{(k)}$ for $y_t \in (H_{k-1}, H_k]$. Then, the objective in (A.10) is concave in y_t with slope $\mathsf{E}_t^Q u_{t+1}^{(k)} - f_t$ for $y_t \in (H_{k-1}, H_k]$.

Let $x_t \in (H_{k-1}, H_k]$, for some $k \in \{2, \dots, T\}$. Consider three cases:

(i) If the slope $\mathsf{E}_t^Q u_{t+1}^{(k-1)} - f_t \leq 0$ (i.e., $V_{tk}^w \leq 0$), then the objective in (A.10) is non-increasing for $y_t \geq H_{k-2}$. Thus, it is optimal to sell $|\underline{C}|$. We have $V_t(x_t, \mathbf{f}_t) = |\underline{C}| f_t + \mathsf{E}_t^Q\big[V_{t+1}(x_t - |\underline{C}|, \mathbf{f}_{t+1})\big]$, which is linear in x_t with slope $\mathsf{E}_t^Q u_{t+1}^{(k-1)}$ for $x_t \in (H_{k-1}, H_k]$.

(ii) If the slopes $\mathsf{E}_t^Q u_{t+1}^{(k)} - f_t \leq 0$ and $\mathsf{E}_t^Q u_{t+1}^{(k-1)} - f_t > 0$ (i.e., $V_{tk}^a \geq 0$ and $V_{tk}^w > 0$), then the objective in (A.10) is increasing in y_t for $y_t \leq H_{k-1}$ and non-increasing for $y_t \geq H_{k-1}$. The optimal decision is $y_t^* = H_{k-1}$; the value function is $V_t(x_t, \mathbf{f}_t) = (x_t - H_{k-1}) f_t + \mathsf{E}_t^Q\big[V_{t+1}(H_{k-1}, \mathbf{f}_{t+1})\big]$, which is linear in x_t with slope f_t for $x_t \in (H_{k-1}, H_k]$.

(iii) If the slope $\mathsf{E}_t^Q u_{t+1}^{(k)} - f_t > 0$ (i.e., $V_{tk}^a < 0$), then the objective in (A.10) is increasing in y_t for $y_t \leq x_t$, and the optimal decision is $y_t^* = x_t$. Under the optimal decision, $V_t(x_t, \mathbf{f}_t) = \mathsf{E}_t^Q\big[V_{t+1}(x_t, \mathbf{f}_{t+1})\big]$ and has slope $\mathsf{E}_t^Q u_{t+1}^{(k)}$ for $x_t \in (H_{k-1}, H_k]$.

In sum, for $x_t \in (H_{k-1}, H_k]$, $k \geq 2$, $V_t(x_t, \mathbf{f}_t)$ is linear in x_t with slope:

$$\begin{cases} \mathsf{E}_t^Q u_{t+1}^{(k-1)}, & \text{if } f_t \geq \mathsf{E}_t^Q u_{t+1}^{(k-1)}, \\ f_t, & \text{if } \mathsf{E}_t^Q u_{t+1}^{(k)} \leq f_t < \mathsf{E}_t^Q u_{t+1}^{(k-1)}, \\ \mathsf{E}_t^Q u_{t+1}^{(k)}, & \text{if } f_t < \mathsf{E}_t^Q u_{t+1}^{(k)}, \end{cases}$$

which is essentially $u_t^{(k)} = k$-th largest element of $\big\{f_t,\, \mathsf{E}_t^Q \mathbf{u}_{t+1}\big\}$.

Finally, when $x_t \in (0, H_1]$, case (iii) above still applies, whereas cases (i) and (ii) are replaced by the following: If $\mathsf{E}_t^Q u_{t+1}^{(1)} \leq f_t$ (i.e., $V_{t1}^a \geq 0$), then the optimal decision is $y_t^* = 0$; the value function is $V_t(x_t, \mathbf{f}_t) = x_t f_t + \mathsf{E}_t^Q\big[V_{t+1}(H_{k-1}, \mathbf{f}_{t+1})\big]$, which is linear in x_t with slope f_t for $x_t \in (0, H_1]$. This, together with case (i), implies that $u_t^{(1)} = \max\big\{f_t,\ \mathsf{E}_t^Q \mathbf{u}_{t+1}\big\}$. ∎

Proof of Proposition 5. The N-period problem is as follows:

$$V_t(x_t, \mathbf{f}_t) = \max_{y_t \in [\underline{y}(x_t), \overline{y}(x_t)]} \ r(y_t - x_t, f_t) + \mathsf{E}_t^Q\big[V_{t+1}(y_t, \mathbf{f}_{t+1})\big], \tag{A.12}$$

$$V_N(x_N, \mathbf{f}_N) = -f_N\, \underline{\lambda}(x_N). \tag{A.13}$$

We inductively prove that for any $\mathbf{f}_t$, $V_t(x_t, \mathbf{f}_t)$ is a concave piece-wise linear function in x_t with slope $v_t^{(k)}$ defined in (24) for $x_t \in (H_{k-1}, H_k]$, $k = 1, \ldots, T$. This is true for $t = N$, as seen in the proof for Proposition 4. Suppose $V_{t+1}(y_t, \mathbf{f}_{t+1})$ is concave in y_t with slope $v_{t+1}^{(k)}$ for $y_t \in (H_{k-1}, H_k]$. Then, the objective in (A.12) is concave in y_t with slope $\mathsf{E}_t^Q v_{t+1}^{(k)} - f_t$ for $y_t \in (H_{k-1}, H_k]$.

Let $x_t \in (H_{k-1}, H_k]$, for some $k \in \{2, \ldots, T-1\}$. Consider five cases below. The first two cases parallel those in the proof of Proposition 4.

(i) If the slope $\mathsf{E}_t^Q v_{t+1}^{(k-1)} - f_t \le 0$ (i.e., $V_{tk}^w \le V_{t,k-1}^c$), it is optimal to sell $|\underline{C}|$. The value function $V_t(x_t, \mathbf{f}_t)$ has slope $\mathsf{E}_t^Q v_{t+1}^{(k-1)}$ for $x_t \in (H_{k-1}, H_k]$.

(ii) If the slopes $\mathsf{E}_t^Q v_{t+1}^{(k)} - f_t \le 0$ and $\mathsf{E}_t^Q v_{t+1}^{(k-1)} - f_t > 0$ (i.e., $V_{tk}^a + V_{tk}^c \ge 0$ and $V_{tk}^w > V_{t,k-1}^c$), the optimal decision is $y_t^* = H_{k-1}$. The value function $V_t(x_t, \mathbf{f}_t)$ has slope f_t for $x_t \in (H_{k-1}, H_k]$.

(iii) If the slopes $\mathsf{E}_t^Q v_{t+1}^{(k)} - f_t > 0$ and $\mathsf{E}_t^Q v_{t+1}^{(k)} - f_t^b \le 0$ (i.e., $V_{tk}^a + V_{tk}^c < 0 \le V_{tk}^{ab} + V_{tk}^c$), then the objective in (A.12) is increasing in y_t for $y_t \le x_t$, and non-increasing for $y_t \ge x_t$. The optimal decision is $y_t^* = x_t$, and $V_t(x_t, \mathbf{f}_t) = \mathsf{E}_t^Q\big[V_{t+1}(x_t, \mathbf{f}_{t+1})\big]$ has slope $\mathsf{E}_t^Q v_{t+1}^{(k)}$ for $x_t \in (H_{k-1}, H_k]$.

(iv) If the slopes $\mathsf{E}_t^Q v_{t+1}^{(k)} - f_t^b > 0$ and $\mathsf{E}_t^Q v_{t+1}^{(k+1)} - f_t^b \le 0$ (i.e., $V_{tk}^{ab} + V_{tk}^c < 0 \le V_{t,k+1}^{ab} + V_{t,k+1}^c$), then the objective in (A.12) is increasing in y_t for $y_t \le H_k$, and non-increasing for $y_t \ge H_k$. The optimal decision is to buy up to $y_t^* = H_k$, and $V_t(x_t, \mathbf{f}_t) = -(H_k - x_t) f_t^b + \mathsf{E}_t^Q\big[V_{t+1}(H_k, \mathbf{f}_{t+1})\big]$ has slope f_t^b for $x_t \in (H_{k-1}, H_k]$.

(v) If the slope $\mathsf{E}_t^Q v_{t+1}^{(k+1)} - f_t^b > 0$ (i.e., $V_{t,k+1}^{ab} + V_{t,k+1}^c < 0$), then the objective in (A.12) is increasing for $y_t \le H_{k+1}$. It is optimal to buy $\overline{C}$, and the resulting value function $V_t(x_t, \mathbf{f}_t) = -\overline{C} f_t^b + \mathsf{E}_t^Q\big[V_{t+1}(x_t + \overline{C}, \mathbf{f}_{t+1})\big]$ has slope $\mathsf{E}_t^Q v_{t+1}^{(k+1)}$ for $x_t \in (H_{k-1}, H_k]$.

In sum, for $x_t \in (H_{k-1}, H_k]$, $k \ge 2$, $V_t(x_t, \mathbf{f}_t)$ is linear in x_t with slope:

$$\begin{cases} \mathsf{E}_t^Q v_{t+1}^{(k-1)}, & \text{if } f_t \ge \mathsf{E}_t^Q v_{t+1}^{(k-1)}, \\ f_t, & \text{if } \mathsf{E}_t^Q v_{t+1}^{(k)} \le f_t < \mathsf{E}_t^Q v_{t+1}^{(k-1)}, \\ \mathsf{E}_t^Q v_{t+1}^{(k)}, & \text{if } f_t < \mathsf{E}_t^Q v_{t+1}^{(k)} \le f_t^b, \\ f_t^b & \text{if } \mathsf{E}_t^Q v_{t+1}^{(k+1)} \le f_t^b < \mathsf{E}_t^Q v_{t+1}^{(k)}, \\ \mathsf{E}_t^Q v_{t+1}^{(k+1)} & \text{if } f_t^b < \mathsf{E}_t^Q v_{t+1}^{(k+1)}, \end{cases}$$

which is essentially $v_t^{(k)} = (k+1)$-th largest element of $\big\{f_t,\ f_t^b,\ \mathsf{E}_t^Q \mathbf{v}_{t+1}\big\}$, When $x_t \in (0, H_1]$, cases (i) and (ii) are replaced by the following: If the slope $\mathsf{E}_t^Q v_{t+1}^{(1} - f_t \leq 0$ (i.e., $V_{t1}^a + V_{t1}^c \geq 0$), we have $y_t^* = 0$, and $V_t(x_t, \mathbf{f}_t)$ has slope f_t for $x_t \in (0, H_1]$.

When $x_t \in (H_{T-1}, K]$, cases (iv) and (v) are replaced by the following: If the slope $\mathsf{E}_t^Q v_{t+1}^{(T)} - f_t^b > 0$ (i.e., $V_{tT}^{ab} + V_{tT}^c < 0$), we have $y_t^* = K$, and $V_t(x_t, \mathbf{f}_t)$ has slope f_t^b for $x_t \in (H_{T-1}, K]$. ∎

D. Lower Bounds on the Value Loss from RI Policy

In this section, we show that if $f_{11} \geq \max\{f_{12}, f_{13}\}$ and $V^w > V^c$, then the expected loss of the RI policy is at least $(V^w - V^c)(H - \underline{y}(x_1))$. If $f_{12} < f_{11} < \min\{f_{12}^b, f_{13}\}$ and $V^w < V^c$, then the expected loss of the RI policy is at least $(V^c - V^w)(H - \max\{\underline{y}(x_1), H'\})$.

In the appendix of the paper, the derivation of the objective (11) suggests that:

$$V(x_1, \mathbf{f}_1) = \max_{y_1 \in [\underline{y}(x_1), H]} U_1(x_1, y_1, \mathbf{f}_1) \equiv V^w y_1 + V^c \min\{H - y_1, \overline{\lambda}(y_1)\} + f_{11} x_1.$$

Proposition 1 shows that the RI policy is optimal for the last two periods. Hence,

$$V_1(x_1, \mathbf{f}_1) - V_1^{\mathrm{RI}}(x_1, \mathbf{f}_1) = U_1(x_1, y_1^*, \mathbf{f}_1) - U_1(x_1, y_1^\dagger, \mathbf{f}_1).$$

We now prove the two statements in sequence.

(i) When $f_{11} \geq \max\{f_{12}, f_{13}\}$ and $V^w > V^c$, Lemma 1(b) and Lemma 2(a) imply that $y_1^\dagger = \underline{y}(x_1) < H = y_1^*$. Then,

$$\begin{aligned} U_1(x_1, y_1^*, \mathbf{f}_1) - U_1(x_1, y_1^\dagger, \mathbf{f}_1) &= V^w H - V^w y_1^\dagger - V^c \min\{H - y_1^\dagger, \overline{\lambda}(y_1^\dagger)\} \\ &\geq V^w(H - y_1^\dagger) - V^c(H - y_1^\dagger) \\ &= (V^w - V^c)(H - \underline{y}(x_1)). \end{aligned}$$

(ii) When $f_{12} < f_{11} < \min\{f_{12}^b, f_{13}\}$, Lemma 2(b) implies that $y_1^\dagger = H$. When $V^w < V^c$, the optimal solution is determined by Lemma 1(a) or (c).

If $\underline{y}(x_1) \geq H'$, then $y_1^* = \underline{y}(x_1)$ and

$$U_1(x_1, y_1^*, \mathbf{f}_1) - U_1(x_1, y_1^\dagger, \mathbf{f}_1) = V^w \underline{y}(x_1) + V^c(H - \underline{y}(x_1)) - V^w H = (V^c - V^w)(H - \underline{y}(x_1)).$$

If $\underline{y}(x_1) < H'$, then $y^* \in [\underline{y}(x_1), H']$ and

$$\begin{aligned} U_1(x_1, y_1^*, \mathbf{f}_1) - U_1(x_1, y_1^\dagger, \mathbf{f}_1) &\geq U_1(x_1, H', \mathbf{f}_1) - U_1(x_1, H, \mathbf{f}_1) \\ &\geq V^w H' + V^c(H - H') - V^w H = (V^c - V^w)(H - H'). \end{aligned}$$

Summarizing the above two cases, we have

$$U_1(x_1, y_1^*, \mathbf{f}_1) - U_1(x_1, y_1^\dagger, \mathbf{f}_1) \geq (V^c - V^w)(H - \max\{\underline{y}(x_1), H'\}).$$