

Merchant Commodity Storage and Term Structure Model Error

Nicola Secomandi,¹ Guoming Lai,² François Margot,¹ Alan Scheller-Wolf,¹ Duane J. Seppi¹

¹Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA
15213-3890, USA

²McCombs School of Business, University of Texas at Austin, 1 University Station, B6000,
GSB 3.136, Austin, TX 78712-1178, USA

ns7@andrew.cmu.edu, guoming.lai@mcombs.utexas.edu,
{fmargot, awolf, ds64}@andrew.cmu.edu

Tepper Working Paper 2010-E71

September 2010; Revised: May 2011, October 2011, September 2012

Abstract

Merchants operations involves valuing and hedging the cash flows of commodity and energy conversion assets as real options based on stochastic models that inevitably embed model error. In this paper we quantify how empirically calibrated model errors about the futures price term structure affect the valuation and hedging of commodity storage assets, specifically the storage of natural gas, an important energy source. We also explore ways to mitigate the impact of these errors. Our analysis demonstrates the differential impact of term structure model error on natural gas storage valuation versus hedging. We also propose an effective approach to deal with the negative effect of such model error on factor hedging, a specific hedging approach. More generally, our work suggests managerial principles for option valuation and hedging in the presence of term structure model error. These principles should have relevance for the merchant management of other commodity conversion assets and for the management of financial options that also depend on term structure dynamics.

1. Introduction

Merchants manage commodity and energy conversion assets – copper mines, natural gas storage facilities and pipelines, power plants, and oil wells and refineries – as real options on commodity and energy prices (Smith and McCardle 1999, Clewlow and Strickland 2000, Maragos 2002, Eydeland and Wolyniec 2003, Geman 2005, Leppard 2005, Secomandi 2010a, Secomandi and Wang 2012, Wu et al. 2012). This is a complex multi-functional task faced everyday by practitioners.

The key components of effective merchant commodity operations are a valuation model, an operating policy, and risk sensitivities (e.g., deltas) for dynamic hedging. A canonical *merchant operations* problem involves valuing and buying commodity conversion assets (or leasing their capacity) and then using an operating policy to maximize their real option value by trading the underlying physical input and output commodities. In particular, merchants acquire such assets when they appear to be priced at a discount relative to their optimized value. Delta hedging the resulting commodity price exposure with futures contracts is critical because it reduces risk capital charges when funding merchant trading activities. That is, delta hedging can mitigate the costs of financial distress due to mismatches between *financing* payments and *operational* cash flows in the presence of credit frictions (see, e.g., Tirole 2006, §5.4 and references therein), and is analogous to how financial institutions hedge similar risks from their financial derivative trading operations (Hull 2010). Online Appendix A provides a simple example.

Valuation, hedging, and value-maximizing operating policies for commodity real options all rely on assumed models of the risk neutral dynamics of the underlying state variables (Smith and McCardle 1999, Birge 2000, Seppi 2002). In practice, the risk neutral dynamics are not known with certainty. The result is model error. Model error can cause two things to go wrong in commodity merchant operations. First, *valuation model error* can cause merchants to misvalue commodity conversion assets; thereby leading to errors in the prices at which they seek to acquire such assets. Second, *hedging model error* can cause the future hedging payoffs to deviate from the future operational cash flows, thus potentially disrupting the capital risk charge mitigation function of financial hedging. Consequently, the impact of model errors on valuation and hedging performance is critical for merchant commodity operations.

In this paper we focus on commodity storage assets that allow merchants to trade a commodity at different prices over time. Specifically, we consider natural gas storage (Maragos 2002). Natural gas is an important commodity, representing 25% of total U.S. energy consumption (EIA 2010), and is projected to become even more important in the future (EIA

2011a). Natural gas storage capacity amounts to roughly 17% of annual U.S. natural gas demand (EIA 2011b,c). At a current price for natural gas of \$2.5/mmBtu, the value of natural gas peak inventory is about \$10B.

As described in the practice-based literature, merchants use multifactor models for the term structure of futures prices to manage natural gas storage (Clewlow and Strickland 2000, Chapter 8, Maragos 2002, pp. 440 and 449-453, Seppi 2002, Eydeland and Wolyniec 2003, Chapter 5 and pp. 351-367, Gray and Khandelwal 2004a,b, Geman 2005, Chapter 3). We distinguish between two approaches. The first, described by Eydeland and Wolyniec (2003, pp. 351-367) and Gray and Khandelwal (2004a,b), models the full term structure of the futures curve using as many factors as there are futures delivery dates over the term of a storage contract. This approach is related to the string and BGM models used to value fixed income options (Kennedy 1994, Brace et al. 1997, Longstaff et al. 2001), and is analogous to the martingale model of demand forecast evolution (Graves et al. 1986, Heath and Jackson 1994). An alternative approach restricts attention to a few large common factors that dominate the statistical dynamics of the futures curve. A typical number of factors is three, which have qualitative interpretations of level (shifting), slope (tilting), and curvature (bending) effects (Cortazar and Schwartz 1994, Borovkova and Geman 2008, Clewlow and Strickland 2000, Chapter 8, but see Borovkova and Geman 2006 and Bjerksund et al. 2011 for other numbers of factors).

This disagreement on the appropriate specification of the natural gas futures curve model is the motivation for our analysis of term structure model error. The goal of our analysis is twofold: (i) To quantify the numerical impact of empirically calibrated model errors and (ii) to devise strategies to mitigate the impact of model error. Our goal is not to identify the correct futures price model, which is the subject of a separate literature (see, e.g., Routledge et al. 2000, Casassus and Collin-Dufresne 2005, and references therein), but rather to study how empirically realistic amounts of model error in misspecified futures models affect valuation and hedging.

There is an infinite variety of possible model errors, so we look to empirical data to impose realistic structure on the type of model errors we consider. In other words, we study the impact of empirically *likely* model errors. We calibrate our model errors using a flexible family of seasonal multifactor futures curve models (Cortazar and Schwartz 1994, Blanco et al. 2002) that nests both full and lower dimensional specifications. We estimate our price models on New York Mercantile Exchange (NYMEX) natural gas futures price data from 1997 through 2006. We find that three to four factors explain about 99% of the monthly price variance. We then use

our empirical model to investigate the impact of three types of realistic term structure model errors: *Model class error*, in which the assumed risk neutral dynamics of the futures curve used for valuation and hedging do not agree qualitatively with the correct risk neutral model; *factor cardinality error*, in which the assumed number of factors used for valuation and hedging is incorrect; and *parameter error*, in which the fitted factor loading coefficients are misestimated.

The choice of the operational trading policy directly affects the valuation of a commodity conversion asset and, thus, the delta hedges corresponding to that valuation. We study natural gas storage operated based on a rolling intrinsic policy, because computing an optimal inventory policy is impractical, due to the curse of dimensionality, with high-dimensional futures price models, and because this heuristic policy is both near optimal and commonly used in practice (see Lai et al. 2010 and references therein, Thompson 2012, and Wu et al. 2012).

There are several approaches to delta hedge option cash flows (Hull 2010, Chapter 6) depending on which futures contracts are used to do the hedging. All of these delta hedging strategies are equivalent in the absence of model error, but they may perform differently in the presence of model error. *Bucket hedging* trades futures contracts with delivery dates corresponding to each date on which operational trading has a cash flow, irrespective of the number of factors in the futures price model used for valuation. The bucket hedging positions are equal to the deltas of the storage option corresponding to these futures prices. In contrast, *factor hedging* involves trading a number of futures contracts equal to the assumed number of random factors in the futures price model (Cortazar and Schwartz 1994, Schwartz 1997, Clewlow and Strickland 2000, §9.5). The parsimonious nature of the low dimensional factor hedging approach is convenient in practice because it involves trading a small number of futures contracts. Moreover, in contrast to bucket hedging, factor hedging remains viable when the operational cash flow dates do not coincide with liquid futures delivery dates.

We show that the position in any individual futures contract in a factor hedging strategy is the sum of (i) the delta of the storage option with respect to that futures price and (ii) a linear combination of additional deltas. We extend the approach of Broadie and Glasserman (1996), which does not include inventory decisions, to derive unbiased pathwise delta expressions for a feasible inventory policy, under a Lipschitz continuity and derivative characterization assumption on its value function, and show that these expressions hold for an optimal inventory policy, under a mild assumption on the operational parameters of our model. The bias resulting from using our expressions with a heuristic policy that does not satisfy our conditions for delta unbiasedness should be small so long as such a policy is near optimal. Our numerical results,

which are enabled by the application of our expressions, suggest that there is little, if any, bias in the delta estimates obtained with the rolling intrinsic policy.

We first study the impact of factor cardinality error (due to omitted factors) and model class error (due to omitted idiosyncratic noise) on natural gas storage valuation. We find that, in practice, these model errors do not appear to be empirically important for valuation purposes, and can be effectively dealt with by including the most important statistical futures price factors in the valuation model. This finding is expected given our empirical calibration results, but sets the stage for our hedging analysis.

We next study the impact of model error on delta hedging. We examine bucket hedging and two versions of factor hedging, testing them using both simulated futures price paths and historical futures prices. We find that bucket hedging is remarkably robust to empirically calibrated model error. In simulated data, it performs near optimally, relative to perfect hedging, in the presence of factor cardinality error (omitted factors) and model class error (omitted idiosyncratic noise). In contrast, factor hedging with model error is very sensitive to which futures contracts are used to implement the hedges. For example, *naïve factor hedging*, which takes positions in futures contracts whose maturities are close together, performs well without model error, but its performance is mostly disastrous in the presence of model error due to omitted factors or omitted idiosyncratic noise: In sharp contrast to our valuation results, ignoring factors that explain less than 1% of the observed price variance (whether factor-driven or idiosyncratic) can lead to surprisingly catastrophic hedging performance. We obtain similar results for historical data, which arguably reflects the combined effects of model class errors, factor cardinality errors, and parameter errors.

We observe that, and explain why, the positions in naïve factor hedging can be extremely large; hence, even statistically small amounts of omitted randomness can be magnified to produce large hedging errors. This finding suggests that in the presence of model error a factor hedging method should choose futures contracts by minimizing the size of their trading positions, according to some norm. In general, this gives rise to a constrained optimization problem, to be solved repeatedly over time. Although in practice it may be practical to solve such a problem, this is numerically impractical in our large-scale simulation study. We thus develop *fine-tuned factor hedging*, which efficiently chooses futures contracts to heuristically minimize the sum of the size (absolute value) of their positions, i.e., we use the 1-norm. This yields trading positions that are much smaller than in naïve factor hedging. Consequently, small price model errors are not magnified to the same extent. In fact, both on simulated and real data,

fine-tuned factor hedging achieves near optimal performance, comparable to bucket hedging but trading many fewer contracts, when the most important statistical factors are included.

Our analysis improves our understanding of the differential effect of realistic term structure model errors on the merchant management of natural gas storage, and provides an effective method to mitigate the negative effect of model error on factor hedging. Although illustrated for natural gas storage, our research suggests some general managerial principles on how to use multifactor models to value and hedge other commodity conversion assets and financial options that depend on other underlying term structures in the presence of model error, e.g., commodity processing and refining assets (Boyabatli et al. 2011, Devalkar et al. 2011, Wu and Chen 2010), commodity swing and put-call Bermudan options (Jaillet et al. 2004, Detemple 2006), and mortgages and interest rate caps and floors (Veronesi 2010). Our work also may inform current commodity storage practice.

The rest of this paper is organized as follows. We review the relevant literature in §2. We present the family of futures price models used in our analysis in §3. We discuss the valuation and hedging of natural gas storage in §4 and §5, respectively. We conduct our model error analysis in §6. We conclude in §7. Online Appendix B extends some of the analysis of §5 to the case of optimal storage valuation. Online Appendix C includes proofs.

2. Literature Review

Previous research on multifactor models of commodity prices focuses on the ability of different models to explain the statistical variability in price time-series data (Gibson and Schwartz 1990, Cortazar and Schwartz 1994, Schwartz 1997, Schwartz and Smith 2000, Blanco et al. 2002, Tolmasky and Hindanov 2002, Casassus and Collin-Dufresne 2005, Geman and Nguyen 2005, Borovkova and Geman 2006, and Suenaga et al. 2008). In contrast, we focus on real option valuation and hedging performance for a family of realistically misspecified multifactor models.

Other authors have studied the valuation of natural gas storage (Chen and Forsyth 2007, Boogert and de Jong 2008, Thompson et al. 2009, Carmona and Ludkovski 2010, Lai et al. 2010, Secomandi 2010b, Thompson 2012, Wu et al. 2012), but they do not consider valuation and hedging with term structure model error. Bjerksund et al. (2011) and Boogert and de Jong (2011/12) emphasize the importance of the choice of price model when valuing natural gas storage, but do not carry out the hedging analysis that we perform in this paper.

Longstaff et al. (2001), Fan et al. (2001), and Driessen et al. (2003) illustrate the importance of using multifactor models to value and hedge interest rate caps and swaptions. In particular, Driessen et al. (2003) show that bucket hedging outperforms factor hedging. We confirm this result in natural gas storage, which features inventory, a component notably absent in fixed income derivatives. Moreover, while caps and swaptions are marked-to-market, natural gas storage is not. Thus, our analysis focuses on the ability of different hedging strategies to replicate cash flows from physical trading, rather than on hedging short-term mark-to-market price changes.

The pathwise approach that we extend for delta estimation (Fu and Hu 1995, Broadie and Glasserman 1996, Glasserman 2004, §7.2) is not common for real options. An exception is Secomandi and Wang (2012) who focus on the valuation and hedging of the real option to transport natural gas, which, different from our model, does not include inventory. Our analysis provides sufficient conditions for pathwise deltas to be unbiased with inventory decisions.

Our approach to compute factor hedging positions is related to the factor hedging implementation in Driessen et al. (2003), but differs from the numerical perturbation approaches discussed by Cortazar and Schwartz (1994), Schwartz (1997), Clewlow and Strickland (2000, §9.5), Longstaff et al. (2001), and Fan et al. (2001). In addition, we propose a method to improve the selection of the contracts used to implement factor hedging, whereas in Longstaff et al. (2001), Fan et al. (2001), and Driessen et al. (2003) they are determined exogenously.

Among others, Birge (2000), Van Mieghem (2003), Gaur and Seshadri (2005), Caldentey and Haugh (2006), and Ding et al. (2007) investigate the relevance of financial markets and hedging for capacity and inventory/production management. In contrast, we study the effect of term structure model error on the merchant management of commodity storage assets.

3. A Family of Futures Price Evolution Models

In this section we present a family of multifactor models to describe the risk-neutral dynamics of the first $N > 1$ contracts in the natural gas futures curve over the term $[T_0 := 0, T_{N-1}]$ of a storage contract. We denote the futures price at time $t \in [0, T_m]$ with maturity date T_m , $\forall m \in \mathcal{N} := \{0, \dots, N-1\}$, by $F(t, T_m)$. The spot price at time T_m is $F(T_m, T_m)$. For a fixed number of factors $K \in \{1, \dots, N-1\}$, the risk-neutral dynamics are described by the following stochastic differential equations (Cortazar and Schwartz 1994):

$$\frac{dF(t, T_m)}{F(t, T_m)} = \sum_{k \in \mathcal{K}} \sigma_{m,k}(t) dZ_k(t), \quad \forall m \in \mathcal{N} \setminus \{0\}, t \in [0, T_m], \quad (1)$$

where $\mathcal{K} := \{0, \dots, K-1\}$ is the set of factor labels, $\sigma_{m,k}(t)$ is the loading coefficient for factor k of the futures price with maturity date T_m at time t , and $dZ_k(t)$ is a standard Brownian motion increment corresponding to factor k . The Brownian motion increments for the different factors all have zero instantaneous correlations. Seasonality in the natural gas price levels is captured by the initial futures prices themselves. This model also allows for possible seasonality in the futures return variance-covariance matrix, because the factor loading coefficients can depend on trading time.

Factor models of futures term structure dynamics are mathematically equivalent to factor models of spot price dynamics (Cortazar and Schwartz 1994). For example, the two-factor spot price model of Schwartz and Smith (2000) and Pilipovic (2007, §5.3.2.2, §6.7) is equivalent to a two-factor futures curve model with parallel shift loadings on one factor and decaying loadings on a second factor (see Ross 1997 and Jaillet et al. 2004 for related models). We conduct our analysis using a dynamic futures curve specification both because of the empirical flexibility of futures curve models and because, in practice, futures are widely used for hedging (Clewlow and Strickland 2000, §8.5, Chapter 9).

We implement our futures price model using an empirical specification that is flexible enough to capture important monthly seasonalities in natural gas price dynamics. It is also consistent with industry practice (Clewlow and Strickland 2000, §8.6, Blanco et al. 2002). Each coefficient $\sigma_{m,k}(\cdot)$ is restricted to be a constant $\sigma_{m,k,n}$ over each time interval $[T_n, T_{n+1})$, with $n \in \mathcal{N}$ and $m \in \mathcal{N}_{n+1} := \{n+1, \dots, N-1\}$. Under this specification, given n and $m \in \mathcal{N}_{n+1}$, $t \in [T_n, T_{n+1})$, and $t' \in (T_n, T_{n+1}]$, with $t' > t$, the futures price $F(t', T_m)$ conditional on $F(t, T_m)$ is

$$F(t', T_m) = F(t, T_m) \exp \left[-\frac{1}{2}(t' - t) \sum_{k \in \mathcal{K}} \sigma_{m,k,n}^2 + \sqrt{t' - t} \sum_{k \in \mathcal{K}} \sigma_{m,k,n} Y_k \right], \quad (2)$$

where $\mathbf{Y} := (Y_k, k \in \mathcal{K})$ is a vector of K independent standard normals. In the special case when $K = N - 1$ and the factor loading coefficients $\sigma_{m,k,n}$ are constant across different time intervals, model (1) reduces to the multifactor version of the homoschedastic Black (1976) model discussed in the practice-based literature (Eydeland and Wolyniec 2003, Chapter 5, Gray and Khandelwal 2004a, Gray and Khandelwal 2004b), and used by Lai et al. (2010) to value natural gas storage.

Our empirical calibration of model (1) is based on the factor loading coefficients being the same under the risk neutral and objective futures curve dynamics. Following industry practice again (Clewlow and Strickland 2000, §8.6, Blanco et al. 2002), we use a principal component analysis (PCA) to estimate the factor loadings for each month separately, as discussed in §6.

4. Valuation

This section is based in part on Lai et al. (2010, §2, §§3.2-3.3). Natural gas storage contracts give merchants the right to inject, store, and withdraw natural gas at a storage facility over a finite time horizon, subject to capacity and inventory constraints. Capacity constraints specify the maximal amount of natural gas, measured in million British thermal units (mmBtus), which a merchant can inject or withdraw per unit of time. Inventory constraints specify the minimal and maximal amounts of natural gas inventory that the merchant can hold at any given point in time. A contract also specifies injections and withdrawal charges, and fuel losses.

We assume that the contracted storage facility is located nearby a liquid wholesale spot market, in which the merchant performs its physical trading. This is realistic for North America (which has roughly one hundred geographically dispersed wholesale markets for natural gas) and for the United Kingdom (the National Balancing Point), Belgium (Zeebrugge), Germany (Emden), and the Netherlands (the Title Transfer Facility). If the storage facility is not located at a market hub, then transport capacity and cost for the pipeline connecting the storage facility to the relevant hub also must be included. We also assume that a natural gas futures market is associated with the wholesale physical market. This is again realistic for North America where NYMEX and ICE trade futures contracts with delivery at Henry Hub, Louisiana and basis swaps for about forty other locations (financially settled forward locational price differences relative to Henry Hub). In the United Kingdom, ICE trades natural futures associated with the National Balancing Point. The European Energy Exchange also trades natural gas futures.

We are interested in valuing the “monthly volatility” component of a storage contract (i.e., the cash flows associated with making physical natural gas trading decisions on a monthly basis). This is realistic because the spot market in the United States is most liquid during the monthly “bid week” in which monthly blocks of gas for the ensuing month are traded.

The timing of the physical trading/inventory decisions is assumed to coincide with the futures maturity dates. The stage set includes the futures maturities. Let a denote an action. A purchase-and-injection is a negative action (i.e., it generates a negative cash flow), a withdrawal-and-sale is a positive action, and the do-nothing action is zero. Natural gas purchased (respectively, sold) at time T_n is available (respectively, unavailable) in storage at time T_{n+1} . A nonzero action thus represents a commodity flow in between two successive stages.

The storage contract minimal and maximal inventory levels are 0 and $\bar{x} \in \mathfrak{R}_+$, respectively, making the set of feasible inventory levels $\mathcal{X} := [0, \bar{x}]$. The injection and withdrawal capacities

per stage are $C^I < 0$ and $C^W > 0$, respectively. The feasible injection action, withdrawal action, and action sets, respectively, with feasible inventory x are $\mathcal{A}^I(x) := [C^I \vee (x - \bar{x}), 0]$, $\mathcal{A}^W(x) := [0, x \wedge C^W]$, and $\mathcal{A}(x) := \mathcal{A}^I(x) \cup \mathcal{A}^W(x)$ ($\cdot \wedge \cdot \equiv \min\{\cdot, \cdot\}$ and $\cdot \vee \cdot \equiv \max\{\cdot, \cdot\}$). We model in-kind fuel losses using the coefficients $\phi^W \in (0, 1]$ and $\phi^I \geq 1$ for withdrawals and injections, respectively. The marginal withdrawal and injection costs are c^W and c^I , respectively. Thus, the per stage cash flow function given an action a and the spot price s is

$$p(a, s) := \begin{cases} (\phi^I s + c^I) a, & a \in \mathfrak{R}_-, \\ 0, & a = 0, \\ (\phi^W s - c^W) a, & a \in \mathfrak{R}_+. \end{cases} \quad \forall s \in \mathfrak{R}_+,$$

The storage contract is managed using a feasible physical trading policy π . Using the more compact notation $F_{n,m} = F(T_n, T_m)$ for futures price, let $\mathbf{F}_n := (F_{n,m}, m \in \mathcal{N}_n)$, $\forall n \in \mathcal{N}$, and $\mathbf{F}'_n := (F_{n,m}, m \in \mathcal{N}_{n+1})$, $\forall n \in \mathcal{N} \setminus \{N-1\}$, denote the futures curves at time T_n inclusive and exclusive, respectively, of the spot price, $s_n \equiv F_{n,n}$, with the conventions $\mathbf{F}_N := 0$ and $\mathbf{F}'_{N-1} := 0$. Let $A_n^\pi(x, \mathbf{F}_n)$ be the decision rule of policy π at stage n in state (x, \mathbf{F}_n) and let x_n^π be the inventory level reached at stage n by policy π . Denote risk neutral expectation by \mathbb{E} and the per stage risk-free discount factor by δ . The value function corresponding to policy π in state (x, \mathbf{F}_n) at stage n is

$$V_n^\pi(x, \mathbf{F}_n) := \sum_{m=n}^{N-1} \delta^{m-n} \mathbb{E}[p(A_m^\pi(x_m^\pi, \mathbf{F}_m), s_m) | x, \mathbf{F}_n]. \quad (3)$$

The optimal storage valuation problem can be formulated as a Markov decision process (MDP; see Appendix B for this formulation). Unfortunately, the MDP for natural gas storage is intractable due to the curse of dimensionality when decisions are conditioned on the entire futures curve \mathbf{F}_n at each stage. Therefore, in our numerical analysis in §6 we assume that the natural gas storage option is managed using the rolling intrinsic policy. (However, our theoretical hedging analysis in §5 is more general and applies to other feasible policies too.) The rolling intrinsic policy is a widely used heuristic among practitioners and has been shown by Lai et al. (2010) to be nearly optimal (see also Thompson 2012, Wu et al. 2012). This policy is based on sequential reoptimization of the intrinsic value model (i.e., the deterministic equivalent of the MDP for exact storage valuation). Its value function can be estimated by a Monte Carlo simulation of the evolution of the futures curve. We refer the reader to Lai et al. (2010, §§3.2-3.3) for details on this policy. We expect that our analysis would carry over to other plausible operating policies.

5. Hedging

In this section, we present methods for hedging the real option to store natural gas. The value function $V_n^\pi(\cdot, \cdot)$ in (3) is defined over a discrete time set. As hedging, in theory, occurs in continuous time, we map this discrete time function into a continuous time value function. On a date $t \in [T_{n-1}, T_n]$ for $n \in \mathcal{N} \setminus \{0\}$, we denote the vector of futures prices as $\mathbf{G}_n(t) := (F(t, T_m), m \in \mathcal{N}_n)$ (at time $t = T_n$ this vector includes the spot price s_n). For every $n \in \mathcal{N} \setminus \{0\}$, $t \in [T_{n-1}, T_n]$, and state $(x, \mathbf{G}_n(t)) \in \mathcal{X} \times \mathfrak{R}_+^{N-n}$, we define the continuous time value function

$$U_n^\pi(t, x, \mathbf{G}_n(t)) := \bar{\delta}(t, T_n) \mathbb{E}[V_n^\pi(x, \mathbf{F}_n) | \mathbf{G}_n(t)], \quad (4)$$

where $\bar{\delta}(t, T_n)$ is the risk-free discount factor from time T_n back to time t ($\bar{\delta}(T_{n-1}, T_n) \equiv \delta$). The function $U_n^\pi(t, x, \mathbf{G}_n(t))$ is the time t updated value of having x inventory units in storage at time T_n given the vector of futures prices at time t when using policy π to manage storage. Starting in state $(x_{n-1}, \mathbf{F}_{n-1})$ at time T_{n-1} , we are interested in hedging changes in the value of $U_n^\pi(t, x_{n-1} - A_{n-1}^\pi(x_{n-1}, \mathbf{F}_{n-1}), \mathbf{G}_n(t))$ due to changes in the futures curve $\mathbf{G}_n(t)$ during time interval $(T_{n-1}, T_n]$.

5.1 Bucket and Factor Hedging

Under a given feasible operating policy π , for every $n \in \mathcal{N} \setminus \{0\}$, we define the delta at time $t \in [T_{n-1}, T_n)$ for maturity T_m , with $m \in \mathcal{N}_n$, given $(x, \mathbf{G}_n(t)) \in \mathcal{X} \times \mathfrak{R}_+^{N-n}$ as

$$\Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) := \frac{\partial U_n^\pi(t, x, \mathbf{G}_n(t))}{\partial F(t, T_m)}. \quad (5)$$

When $t = T_{n-1}$ we are interested in the partial derivative of $U_n^\pi(T_{n-1}, x, \mathbf{G}_n(T_{n-1}))$ with respect to each $F(T_{n-1}, T_m)$ with x held fixed, because at time T_{n-1} the deltas are computed right after the *commercial* implementation of a feasible action at this time. Otherwise, when $t \in (T_{n-1}, T_n)$, inventory is fixed based on the prior decision at T_{n-1} . By Ito's lemma, the dynamics of $U_n^\pi(t, x, \mathbf{G}_n(t))$ in $(T_{n-1}, T_n]$ are

$$dU_n^\pi(t, x, \mathbf{G}_n(t)) = D(t)dt + \sum_{m \in \mathcal{N}_n} \Delta_{n,m}^\pi(x, t, \mathbf{G}_n(t)) dF(t, T_m), \quad (6)$$

with

$$D(t) := \left[\frac{U_n(t, x, \mathbf{G}_n(t))}{\partial t} + \frac{1}{2} \sum_{m \in \mathcal{N}_n} \sum_{l \in \mathcal{N}_n} \frac{\partial^2 U_n(t, x, \mathbf{G}_n(t))}{\partial F(t, T_m) \partial F(t, T_l)} \sum_{k \in \mathcal{K}} \sigma_{m,k,n-1} F(t, T_m) \sigma_{l,k,n-1} F(t, T_l) \right].$$

Denote by $q_{n,m}^\pi(t, x, \mathbf{G}_n(t))$ the replicating position corresponding to the futures contract with maturity T_m , $m \in \mathcal{N}_n$, at time t when hedging $U_n^\pi(t, x, \mathbf{G}_n(t))$. Consider a portfolio that is long physical storage and continuously shorts these futures positions. The dynamics of this portfolio's value $\Pi_n^\pi(t, x, \mathbf{G}_n(t))$ satisfy the stochastic differential equation

$$d\Pi_n^\pi(t, x, \mathbf{G}_n(t)) = dU_n^\pi(t, x, \mathbf{G}_n(t)) - \sum_{m \in \mathcal{N}_n} q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_m). \quad (7)$$

Using (6) and (1), (7) can be rearranged as

$$d\Pi_n^\pi(t, x, \mathbf{G}_n(t)) = D(t)dt + \sum_{k \in \mathcal{K}} \left[\sum_{m \in \mathcal{N}_n} \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1} - \sum_{m \in \mathcal{N}_n} q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1} \right] dZ_k(t). \quad (8)$$

It follows from (8) that the stochastic variability is hedged when

$$\sum_{m \in \mathcal{N}_n} q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1} = \sum_{m \in \mathcal{N}_n} \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1}, \forall k \in \mathcal{K}. \quad (9)$$

This is a system of K linear equations (one for each factor) with $N - n$ unknowns (the futures positions $q_{n,m}^\pi(t, x, \mathbf{G}_n(t))$). When the number of random factors driving the futures curve is equal to or more than the number of futures contracts (i.e., $K \geq N - n$), setting $q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) = \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t))$, $\forall m \in \mathcal{N}_n$, solves the system of linear equations (9). Otherwise (i.e., $K < N - n$), this system of linear equations is underdetermined and multiple delta hedge positions exist. Bucket and factor hedging are different particular solutions.

Bucket Hedging. Bucket hedging sets the replicating positions equal to the deltas, i.e., $q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) = \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t))$, $\forall m \in \mathcal{N}_n$, which provides a solution to the system of linear equations (9). In other words, bucket hedging ignores the factor structure of the futures curve dynamics when determining the replicating positions. Thus, bucket hedging only depends on the futures curve factor structure via the value function $U_n^\pi(t, x, \mathbf{G}_n(t))$ used to compute the deltas $\Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t))$. Because of this dependence, the bucket hedges obtained when using different numbers of factors for valuation may differ.

Factor Hedging. Factor hedging takes positions in as many futures contracts as the number of futures curve factors, K , and sets to zero the positions in the remaining $N - n - K$ futures contracts. However, which K contracts are traded is not unique, as, in principle, any K futures

contracts can be used as long as their factor loadings are linearly independent. We include in set \mathcal{M}_n the maturity labels of the K futures contracts used for hedging, and collect the remaining unused maturity labels in set $\mathcal{L}_n := \mathcal{N}_n \setminus \mathcal{M}_n$. We set $q_{n,l}^\pi(t, x, \mathbf{G}_n(t)) := 0$ for all $l \in \mathcal{L}_n$ in the system of linear equations (9), and determine $q_{n,m}^\pi(t, x, \mathbf{G}_n(t))$ for all $m \in \mathcal{M}_n$ by solving the following resulting system of K linear equations in K unknowns:

$$\sum_{m \in \mathcal{M}_n} q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1} = \sum_{m \in \mathcal{N}_n} \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) F(t, T_m) \sigma_{m,k,n-1}, \forall k \in \mathcal{K}. \quad (10)$$

In contrast to bucket hedging, factor hedging depends on the futures curve factor structure both via the factor loadings, $\sigma_{m,k,n-1}$, and via the deltas, $\Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t))$.

Proposition 1 provides insight into the structure of the factor hedging trading positions: Each of these positions is the sum of the delta corresponding to a traded futures contract and a linear combination of the deltas corresponding to the untraded futures contracts. This proposition uses some additional notation. Given $n \in \mathcal{N}$ such that $N - n > K$, let B_{n-1} and E_{n-1} denote the $K \times K$ and $(N - n - K) \times K$ matrices of factor loading coefficients in model (1) for maturity labels in sets \mathcal{M}_n and \mathcal{L}_n , respectively, as of time T_{n-1} . The inverse of B_{n-1} is $\zeta_{n-1}/|B_{n-1}|$, where ζ_{n-1} is the transpose of the matrix of cofactors of B_{n-1} (if $K = 1$ then the cofactor matrix is just the scalar 1) and $|B_{n-1}|$ is the determinant of B_{n-1} . The element in position $(l - (n + K) + 1, m - n + 1)$ of matrix $E_{n-1}\zeta_{n-1}$ is $\gamma_{l,m,n-1}$, with $l \in \mathcal{L}_n$ and $m \in \mathcal{M}_n$.

Proposition 1 (Factor hedging positions). *Pick $n \in \mathcal{N}$ such that $N - n > K$ and let $t \in [T_{n-1}, T_n)$. Suppose that $|B_{n-1}| \neq 0$. For all traded maturity labels $m \in \mathcal{M}_n$, it holds that*

$$q_{n,m}^\pi(t, x, \mathbf{G}_n(t)) = \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) + \sum_{l \in \mathcal{L}_n} \Delta_{n,l}^\pi(t, x, \mathbf{G}_n(t)) \frac{F(t, T_l)}{F(t, T_m)} \frac{\gamma_{l,m,n-1}}{|B_{n-1}|}. \quad (11)$$

Proposition 1 states that factor hedging involves choosing a subset of K futures contracts at each date $t \in [T_{n-1}, T_n)$ and computing the trading positions by inverting the square submatrix B_{n-1} of their loading coefficients, which is the position generating matrix, at time T_{n-1} . Different implementations of factor hedging use different subsets of futures to delta hedge. We consider two: Naïve factor hedging and fine-tuned factor hedging.

Naïve Factor Hedging. At time $t \in [T_{n-1}, T_n)$ naïve factor hedging simply takes positions in the K shortest maturity contracts (i.e., with maturity labels $n, \dots, n + K - 1$). This approach does not attempt to control the size (absolute value) of its trading positions. Size is irrelevant, in theory, in the absence of model error and if there are no liquidity concerns. Size does become

potentially important with model error, because, as will be shown in §6, large trading positions amplify the effect of any model error in the realized price changes.

Fine-Tuned Factor Hedging. This issue of size suggests that in the presence of model error a factor hedging method should trade futures contracts whose positions are the smallest possible, according to some norm. Determining these contracts requires solving a potentially difficult minimum-norm constrained optimization problem at each trading time t and in each state $(x, \mathbf{G}_n(t))$. While this may be practical in reality, it is impractical in our large-scale simulation study in §6. We thus develop fine-tuned factor hedging, which aims heuristically at obtaining small futures trading positions. Based on Proposition 1, at a given trading time t , this method selects the set of K contracts to trade that leads to the position generating matrix with the largest determinant $|B_{n-1}|$. This method is computationally efficient because the choice of which position generating matrix to use at each trading time can be made once up front at time T_0 , as this choice depends only on the trading time itself, rather than also the state $(x, \mathbf{G}_n(t))$.

Remark 1 (Equivalent hedging approaches). The arguments used so far in discussing bucket and factor hedging generalize to hedging portfolios including between $K+1$ and $N-2$ contracts. In other words, bucket and factor hedging are two polar cases of a family of hedging strategies.

5.2 Computing Deltas

The storage deltas are critical inputs to both bucket and factor hedging. Proposition 2 provides a pathwise expression for the deltas. Proposition 3 bounds them.

Assumption 1 allows us to apply in a recursive manner the method of proof of Broadie and Glasserman (1996) to establish Proposition 2. In particular, the second part of this assumption roughly states that the feasible policy π has constant actions in neighborhoods of differentiability of its value function with respect to each element of the futures curve. As shown in Online Appendix B, an optimal operating policy obtained under a mild assumption, in particular, satisfies the properties stated in Assumption 1.

Assumption 1 (Lipschitz continuity and derivative characterization). (a) In every stage $n \in \mathcal{N}$, for each given inventory $x \in \mathcal{X}$ the function $V_n^\pi(x, \mathbf{F}_n)$ is Lipschitz continuous in the futures curve $\mathbf{F}_n \in \mathfrak{R}_+^{N-n}$; that is, there exists $L_n^\pi(x) \in \mathfrak{R}_+$ such that $|V_n^\pi(x, \mathbf{F}_n^2) - V_n^\pi(x, \mathbf{F}_n^1)| \leq L_n^\pi(x) \sum_{m=n}^N |F_{n,m}^2 - F_{n,m}^1|$, $\forall \mathbf{F}_n^1, \mathbf{F}_n^2 \in \mathfrak{R}_+^{N-n}$. (b) Moreover, at every futures curve $\bar{\mathbf{F}}_n$ where

$V_n^\pi(x, \mathbf{F}_n)$ is differentiable with respect to each element of \mathbf{F}_n the policy π has a unique action denoted by $a_n^\pi(x, \bar{\mathbf{F}}_n)$ and

$$\frac{\partial V_n^\pi(x_n, \mathbf{F}_n)}{\partial F_{n,m}} \Big|_{\mathbf{F}_n = \bar{\mathbf{F}}_n} = \frac{\partial p(a_n^\pi(x_n, \bar{\mathbf{F}}_n), s_n)}{\partial s_n} \Big|_{s_n = \bar{s}_n} + \frac{\partial \delta \mathbb{E}[V_{n+1}^\pi(x_n - a_n^\pi(x_n, \bar{\mathbf{F}}_n), \mathbf{F}_{n+1}) | \mathbf{F}'_n]}{\partial F_{n,m}} \Big|_{\mathbf{F}_n = \bar{\mathbf{F}}_n}, \quad \forall m \in \mathcal{N}_n.$$

We let $1\{\mathcal{E}\}$ be the indicator function of event \mathcal{E} .

Proposition 2 (Pathwise deltas). *Under Assumption 1, for every $n \in \mathcal{N} \setminus \{0\}$ it holds that*

$$\Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) = \frac{\bar{\delta}(t, T_m)}{F(t, T_m)} \mathbb{E}[(\phi^I 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) < 0\} + \phi^W 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) > 0\}) s_m A_m^\pi(x_m^\pi, \mathbf{F}_m) | x_n, \mathbf{G}_n(t)], \quad (12)$$

for all $m \in \mathcal{N}_n$, $x_n \in \mathcal{X}$, $t \in [T_{n-1}, T_n)$, and $\mathbf{G}_n(t) \in \mathfrak{R}_+^{N-n}$.

Expression (12) can be used to estimate the deltas of policy π by Monte Carlo simulation simultaneously with the valuation of this policy. In §6 we use this expression with the rolling intrinsic policy, even though its value function can violate the Lipschitz condition in Assumption 1. This delta estimation approach may thus yield biased estimates, but can be implemented efficiently. Alternatively, we could estimate deltas by resimulation, which is both biased and computationally expensive (Glasserman 2004, §7.1). The main difficulty with resimulation in the storage context is its high computational requirement, which makes it impractical for our extensive multimodel simulation analysis.

Our computational results in §6 suggest that our calculated deltas are accurate. As shown in Online Appendix B, an expression analogous to (12) holds under an optimal policy (under a mild assumption). The work of Lai et al. (2010) suggests that the bias in applying (12) with the rolling intrinsic policy should be small, as they find that this policy is near optimal (see also Thompson 2012, Wu et al. 2012). Our computational results are consistent with this statement.

The term on the right hand side of (12) can be interpreted as being proportional to a “weighted average” of the action taken by a feasible policy in a given stage and state. To see this, notice that the ratio $s_m/F(t, T_m)$ is a nonnegative random variable with mean one (because $\mathbb{E}[s_m | F(t, T_m)] = F(t, T_m)$). We can thus think of this ratio as weighing the amount of commodity traded by policy π in state (x_m^π, \mathbf{F}_m) at stage n , that is, $(\phi^I 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) < 0\} + \phi^W 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) > 0\}) A_m^\pi(x_m^\pi, \mathbf{F}_m)$. Consequently, the expression

$$\mathbb{E} \left[\frac{s_m}{F(t, T_m)} (\phi^I 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) < 0\} + \phi^W 1\{A_m^\pi(x_m^\pi, \mathbf{F}_m) > 0\}) A_m^\pi(x_m^\pi, \mathbf{F}_m) | x_n, \mathbf{G}_n(t) \right]$$

is a weighted expectation of the amount of commodity traded by policy π in this state and stage. The delta $\Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t))$ is proportional to this expectation according to $\bar{\delta}(t, T_m)$.

Proposition 3 (Bounds on deltas). *Consider feasible policy π . Under Assumption 1, for every $n \in \mathcal{N} \setminus \{0\}$ it holds that $\bar{\delta}(t, T_m)\phi^I C^I \leq \Delta_{n,m}^\pi(t, x, \mathbf{G}_n(t)) \leq \bar{\delta}(t, T_m)\phi^W C^W$, for all $m \in \mathcal{N}_n$, $x_n \in \mathcal{X}$, $t \in [T_{n-1}, T_n)$, and $\mathbf{G}_n(t) \in \mathfrak{R}_+^{N-n}$.*

Proposition 3 allows us to relate the size of the rolling intrinsic policy deltas that we compute in §6 to those of a policy that does satisfy Assumption 1, in particular an optimal policy (under a mild assumption satisfied in our numerical analysis).

6. Term Structure Model Error Analysis

In this section, we analyze the impact of term structure model error using both simulated and historical futures price data. Using simulated data, we analyze the effects of factor cardinality error and model class error on valuation and hedging. Using real historical data, where the true model is unknown, we study the combined impact of factor cardinality error, model class error, and parameter error on hedging. Our hedging analysis also includes time discretization error (as we do not hedge continuously) and may contain delta computation error (as our computed deltas may be biased and/or imprecise). As these last two types of errors do not seem to be numerically large and do not play a fundamental role in our analysis, we do not focus on them as additional sources of error.

6.1 Data and Price Model Estimation

We use ten years of NYMEX natural gas futures price data from January 1997 to December 2006 to calibrate the coefficients of the multifactor price model (1). As model error will have different impacts in models with different factor loadings, our empirical estimation allows us to measure the impact of model error given a realistic set of factor loadings.

Our estimation approach to calibrate our futures price model is as follows: On each trading day, we compute the daily changes of the futures prices corresponding to the first twenty-three maturities. We do this because we analyze storage contracts with twenty-three month terms (see §6.2). We segregate these price changes into twelve subsets, each including the daily price changes for all the relevant maturities during a specific month, over the ten year sample. We perform twelve separate PCAs on each of these monthly subsets to obtain twelve sets of monthly factor loading coefficients for price model (1). Thus, for a given maturity we have twelve sets




Figure 1: The percentage of the total variance explained by the first four factors in a monthly PCA analysis over January 1997 to December 2006.

of factor loadings, one for each trading month. The seasonal specification gives our calibrated model substantial flexibility, with the nonseasonal specification being just a special case.

Figure 1 displays the percentages of the total variance in the observed price changes explained by the first four factors for each of the twelve PCAs. The first factor captures roughly 86.0-93.1% of this variance, the second factor 5.2-9.6%, the third factor 0.6-3.4%, and the fourth factor always less than 1%. The first three factors explain at least 99% of the total variance in the observed price changes in each month except February and September. The first four factors explain at least 99% of the total variance in the observed price changes in each month. Thus, our discussion focuses on the loading coefficients of the first four factors.

Figure 2 displays the monthly loading coefficients of all the relevant maturities on the first four factors. For expositional convenience we distinguish between months in the heating season (November-March), panels (a), (c), (e), and (g), and the rest of the year, panels (b), (d), (f), and (h). Typically, the coefficients decrease rapidly, with the factor loading coefficients of a given factor being less than half of those of the previous factor.

The first factor changes the slope of the natural gas futures term structure. In particular, the first factor induces positively correlated shocks along the futures term structure (since the

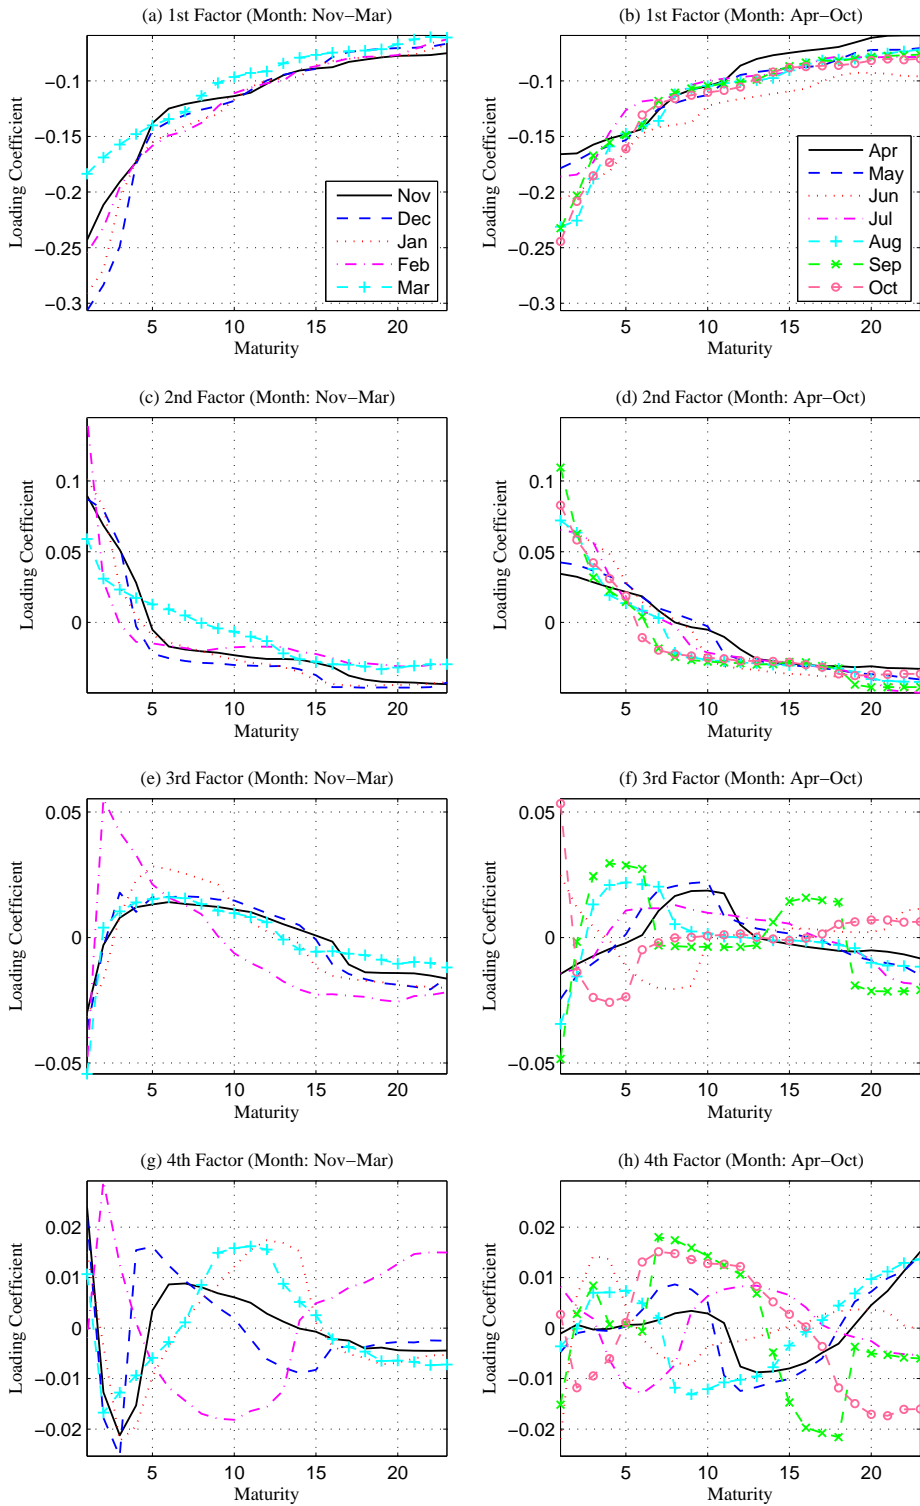


Figure 2: The estimated loading coefficients of the first four factors for maturities one to twenty-three for the price model (1) in each relevant month from a monthly PCA over January 1997 to December 2006.

coefficients are of the same sign), but it moves the short end of the term structure more than the long end (since the coefficients decline in magnitude in time to delivery). The second factor also shocks the slope of the futures term structure, but the changes at the long and short ends of the curve are negatively correlated (since the coefficients change sign). In contrast, the third factor shocks the curvature of the futures term structure (since the signs of short-term and long-term futures coefficients have the opposite sign from the intermediate term coefficients). The fourth factor induces irregular “squiggles” along the futures term structure (since the signs of the coefficients change multiple times). In terms of seasonality, the shocks to the term structure slope induced by the first factor are greater in the heating months than in the rest of the months (since the magnitude of the short-term/long-term coefficient differences are greater). There also appear to be seasonalities in the location and shape of the curvature and squiggle changes.

6.2 Other Parameters

Our experimental setup is related to the twenty-four month valuation analysis of Lai et al. (2010); i.e., $N = 24$. To simplify bookkeeping, we start the storage contract in stage 1 and value it at time 0 (stage 0). Thus, there are operational cash flows at twenty-three monthly stages. As in Lai et al. (2010), we normalize the maximum storage space, \bar{x} , to 1 mmBtu and set the initial inventory, x_1 , equal to 0, the injection and withdrawal capacities, C^I and C^W , respectively, equal to -0.3 and 0.6 mmBtu per month (i.e., 3.33 months to fill up and 1.67 months to empty), the injection and withdrawal marginal costs, c^I and c^W , respectively, equal to $\$0.02/\text{mmBtu}$ and $\$0.01/\text{mmBtu}$, and the injection and withdrawal fuel losses, ϕ^I and ϕ^W , respectively, equal to 1.01 and 0.99. These parameters reflect realistic values for natural gas storage contracts. We use a constant annualized risk-free interest rate equal to 5.05%.

6.3 Valuation Results

We investigate the effect of factor cardinality error and then model class error on valuation.

Factor Cardinality Error. Given our empirical futures price estimation results, we expect the value of storage to increase in the number of considered factors, but to level off once the most important statistical factors have been included. Figure 3 displays the value of storage using the rolling intrinsic policy, $U_1^{RI}(0, 0, \mathbf{F}'_0)$, as a function of the number of factors on four valuation dates representing different seasons of the year: 3/1/2006, 6/1/2006, 8/31/2006, and 12/1/2006. Spring and Winter are part of the heating season. Summer and Fall are not. The displayed values of storage vary in the ranges $\$4.333$ - 4.595 , $\$5.009$ - 5.169 , $\$4.310$ - 4.519 , and




Figure 3: The value of storage as a function of the number of factors in different months.

\$1.217-1.536 in Spring, Summer, Fall, and Winter, respectively. The Monte Carlo standard errors for these valuations are between \$0.012 and \$0.018.

As expected, adding factors to the price model initially increases the value of storage somewhat, but then the valuation flattens out. The initial increases are about 4-5% of the maximal valuation in Spring, Summer, and Fall, and roughly 20% in the Winter. In Spring, Summer, and Fall, three factors are sufficient to achieve valuations that are 98% of the maximum; in Winter, this requires six factors. This is consistent with the earlier observation that four factors explain more than 99% of the observed total daily price change variance.

The lower value of storage in the Winter instance is a consequence of the fact that Winter is the high-price season and the assumption that the initial inventory is zero. Thus, the value of storage due to seasonality, i.e., the intrinsic value, is much lower in the Winter instance than in the other instances. In contrast, the value of storage due to price uncertainty, i.e., the extrinsic value, is more comparable across instances.

We conclude that parsimonious price models can reproduce essentially the same valuations as higher, and in particular full dimensional, price models. To be clear, our point is not that any one of these specifications is correct. Rather, our results suggest that realistic factor cardinality error due to omitted factors in an empirically calibrated version of price model (1) is not problematic for storage valuation, once the most important statistical factors are included.

Model Class Error. Our factor cardinality error results may be affected by the particular correlation structure of the factor models used to generate the simulated data. In practice, however, increasing the number of factors may not bring a price model closer to reality, since the price randomness beyond the fourth factor may simply be idiosyncratic noise, rather than structured correlation. We thus conduct a model class error analysis to investigate the effect of excluding this possibly spurious correlation structure from our factor cardinality error analysis.

To do this, we employ a calibrated price model with four factors and maturity-specific idiosyncratic noise added to the (natural logarithm of the) individual futures prices. These noise terms are normally distributed, mutually independent, with zero means and variances equal to the residual price variances once the four factors are removed. The resulting model has the same variance structure as the full twenty-three factor model but a different correlation structure.

The valuations obtained with this model vary between 99% and 101% of those obtained with at least three, four, and six factors in the Spring and Summer, Fall, and Winter cases, respectively. Hence, our conclusions based on our factor cardinality error analysis do not appear to be affected in an important manner by the correlation structure of the omitted randomness.

6.4 Measuring Hedging Effectiveness and Hedging Analysis Setup

As natural gas storage is not traded in a liquid market, market prices are not available to measure hedging effectiveness. Thus, we focus on how effective futures hedging is in replicating the cash flows generated by physical trading along a futures curve sample path. Given an initial futures curve at time 0, we let ω denote a sample path of N futures curve realizations at times 0 through T_{N-1} . In our simulation analysis, futures price paths are generated from different specifications of price model (1), and the idiosyncratic noise variant thereof. Alternatively, we also use paths from historical futures data. We define the time 0 present value of the cumulative cash flows of the rolling intrinsic policy from time T_1 through time T_{N-1} along path ω as $P(\omega) := \sum_{n \in \mathcal{N} \setminus \{0\}} \delta^n p(a_n^{RI}(x_n, \mathbf{F}_n), s_n; \omega)$, where $a_n^{RI}(x_n, \mathbf{F}_n)$ is the action taken by the rolling intrinsic policy in state (x_n, \mathbf{F}_n) at stage n . (We do not superscript $P(\omega)$, as well as $\psi_n(\cdot; \omega)$ and $\Psi(\omega)$ below, by RI to ease the notation.)

We denote the time T_{n-1} vector of replicating futures positions of a given hedging method when using the rolling intrinsic policy as $\mathbf{q}_n^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1}) := (q_{n,m}^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1}), m \in \mathcal{N}_n)$, where “.” stands for $x_1 \equiv 0$ if $n = 1$ and for $x_{n-1} - a_{n-1}^{RI}(x_{n-1}, \mathbf{F}_{n-1})$ if $n > 1$. The cash flow generated at time T_n along path ω by the futures trading strategy $\mathbf{q}_n^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1})$ is

$\psi_n(\mathbf{q}_n^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1}); \omega) := \sum_{m \in \mathcal{N}_n} q_{n,m}^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1})(F_{n,m} - F_{n-1,m})(\omega)$. We define the time 0 present value of these cumulative cash flows from time T_1 through time T_{N-1} as $\Psi(\omega) := \sum_{n \in \mathcal{N} \setminus \{0\}} \delta^n \psi_n(\mathbf{q}_{n-1}^{RI}(T_{n-1}, \cdot, \mathbf{F}'_{n-1}); \omega)$. These definitions assume monthly rebalancing of the futures positions, but our analysis could accommodate more frequent rebalancing by increasing the granularity of our futures curve samples.

From a merchant's perspective, the reduction in physical cash flow variability due to hedging is a relevant metric of hedging effectiveness. We denote by $\text{VAR}(P|x_1, \mathbf{F}'_0)$ and $\text{VAR}(P - \Psi|x_1, \mathbf{F}'_0)$ the cross-path variances of the total physical cash flows, P , and residual total cash flows after hedging, $P - \Psi$, respectively, given x_1 and \mathbf{F}'_0 . Our hedging effectiveness metric (HEM) is the reduction in $\text{VAR}(P|x_1, \mathbf{F}'_0)$ due to hedging:

$$\text{HEM} := 100 \left[1 - \frac{\text{VAR}(P - \Psi|x_1, \mathbf{F}'_0)}{\text{VAR}(P|x_1, \mathbf{F}'_0)} \right].$$

HEM can be estimated on simulated data. For each pair of number of factors used for valuation and hedging method that we consider, we use Monte Carlo simulation to estimate $\text{VAR}(P|x_1, \mathbf{F}'_0)$ and $\text{VAR}(P - \Psi|x_1, \mathbf{F}'_0)$, and hence HEM. HEM equals 100 with a perfect hedge. In this case, $P(\omega) - \Psi(\omega)$ is equal to $U_1^{RI}(0, 0, \mathbf{F}'_0)$, a constant given x_1 and \mathbf{F}'_0 , for almost every sample path ω .

Empirical backtesting of hedging using historical data involves following a hedging strategy over a series of subintervals, each of which represents one observation in the sample. A complication is that the initial futures curve, \mathbf{F}'_0 , and hence the initial valuation, $U_1^{RI}(0, 0, \mathbf{F}'_0)$, vary across subintervals (i.e., observations). We thus modify HEM using the sample variances of $P - \Psi - U_1^{RI}(0, 0, \mathbf{F}'_0)$ and $P - U_1^{RI}(0, 0, \mathbf{F}'_0)$ to control for the different initial values across observations. We denote by BHEM the resulting backtesting version of HEM.

HEM and BHEM are related to the metric used by Driessen et al. (2003) to measure the ratio of the sample variances of hedging errors and changes in the market price of a particular derivative investment. However, we work with total physical and hedging cash flows rather than hedging errors and market price changes, as natural gas storage is not marked-to-market.

Our hedging analysis setup considers triples consisting of (i) a model that generates the price data, (ii) the number of factors, K , in price model (1) used to compute the value and the deltas of the rolling intrinsic policy, and (iii) the choice of futures contracts to trade according to a given hedging strategy.

6.5 Hedging Results with Simulated Data

We analyze factor cardinality and model class errors by performing a series of *what if* experiments. We take the NYMEX closing futures prices on 6/1/2006 as the initial time 0 futures curve. We rebalance the financial trading positions once every month. We use a Monte Carlo sample size of 30 paths to estimate HEM for each valuation model and hedging method pair that we consider. Rebalancing requires nested simulations to compute deltas, for which we use 10,000 additional Monte Carlo samples. Computing the cash flows associated with a given hedging strategy along one sample path requires about 7 Cpu minutes, depending on the price models and hedging method used, using the g++ version 4.3.0 20080428 (Red Hat 4.3.0-8) compiler on a 64 bits Monarch Empro 4-Way Tower Server with four AMD Opteron 852 2.6GHz processors, each with eight DDR-400 SDRAM of 2 GB running Linux Fedora 11 (the stated Cpu times correspond to using a single processor). To control for sampling error, we use common random numbers when generating price samples across different models that share common factors.

Factor Cardinality Error. To study factor cardinality error we simulate futures curve paths from price model (1) with a given number of factors, which we call the data generation number of factors (DGNF). We also use price model (1) to estimate the deltas of the rolling intrinsic policy but with a potentially different number of factors K . The loading coefficients in both the data generating model and the model used for delta estimation are from our empirical calibration of price model (1). Thus, the magnitudes are empirically realistic.

Our factor cardinality error analysis considers both omitted factors and excess factors. With factor hedging, these errors correspond to underhedging due to omitted factors and overhedging due to excess factors, respectively. We also compare these results with the case of no factor cardinality error. We expect all delta hedging strategies to perform well with no factor cardinality error or with excess factors, but their performance should degrade when there are omitted factors. However, in light of the fact that the first few factors explain most of the total variance in futures price changes, our primary question is quantitative: How much does the hedging performance of different strategies degrade as the omitted factors become statistically small?

Naïve factor hedging. Table 1 shows the results for naïve factor hedging. Without factor cardinality error, naïve factor hedging performs well, as expected, with HEM estimates above 96.18 on the diagonal (i.e., where $K = \text{DGNF}$). The HEM estimates along the diagonal of Table 1 are not equal to 100 percent — due to simulation, discrete rebalancing, and delta computation

Table 1: HEM estimates for naïve factor hedging with factor cardinality error. The emphasized values correspond to no model error. The off-diagonal terms include model error.

DGNF	K							
	1	2	3	5	10	15	20	23
1	<i>99.94</i>	96.00	99.51	99.57	99.72	99.72	99.73	99.73
2	87.82	<i>96.18</i>	99.71	99.45	99.82	99.82	99.82	99.82
3	85.20	-5,505.48	<i>98.97</i>	99.36	99.72	99.69	99.70	99.70
5	82.59	-19,600.60	-1,825.09	<i>99.45</i>	99.62	99.68	99.68	99.68
10	79.42	-47,753.00	-11,979.60	-807.05	<i>99.54</i>	99.59	99.60	99.60
15	80.99	-42,717.10	-11,140.60	-1,695.29	-70.11	<i>99.63</i>	99.64	99.64
20	80.93	-42,643.10	-11,258.80	-1,925.69	-207.81	87.15	<i>99.63</i>	99.63
23	79.88	-45,102.80	-11,809.80	-2,844.02	-425.75	84.06	99.62	<i>99.61</i>

Table 2: The cumulative long and short replicating futures positions taken by naïve factor hedging along a price path averaged across simulated paths. The emphasized values correspond to no model error. The off-diagonal terms include model error.

Average Cumulative Long Positions								
DGNF	K							
	1	2	3	5	10	15	20	23
1	<i>11.13</i>	585.04	547.89	1,461.69	1,094.57	291.49	26.51	26.31
2	10.76	<i>577.97</i>	545.05	1,449.67	1,090.61	291.23	26.48	26.27
3	10.65	578.63	<i>545.74</i>	1,452.29	1,092.50	291.45	26.56	26.35
5	10.60	580.90	549.25	<i>1,455.13</i>	1,093.31	290.83	26.50	26.30
10	10.60	578.56	547.75	1,451.76	<i>1,091.76</i>	290.99	26.55	26.34
15	10.62	578.54	547.86	1,452.10	1,091.89	<i>290.97</i>	26.56	26.34
20	10.60	578.26	547.36	1,450.84	1,091.34	290.80	<i>26.53</i>	26.33
23	10.62	578.57	547.46	1,451.05	1,091.42	290.75	26.57	<i>26.36</i>

Average Cumulative Short Positions								
DGNF	K							
	1	2	3	5	10	15	20	23
1	<i>0.66</i>	613.35	539.26	1,427.33	1,102.80	279.90	14.70	14.49
2	0.67	<i>606.22</i>	536.96	1,415.81	1,099.23	280.10	15.13	14.92
3	0.67	607.13	<i>537.81</i>	1,418.53	1,101.29	280.46	15.34	15.13
5	0.68	609.48	541.37	<i>1,421.39</i>	1,102.14	279.88	15.32	15.11
10	0.66	606.72	539.75	1,417.95	<i>1,100.61</i>	280.04	15.38	15.17
15	0.66	606.58	539.69	1,418.27	1,100.73	<i>280.00</i>	15.36	15.15
20	0.66	606.54	539.41	1,417.40	1,100.04	279.84	<i>15.35</i>	15.15
23	0.67	606.89	539.56	1,417.61	1,100.08	279.79	15.39	<i>15.18</i>

Table 3: HEM estimates for fine-tuned factor hedging with factor cardinality error. The emphasized values correspond to no model error. The off-diagonal entries include model error.

DGNF	K							
	1	2	3	5	10	15	20	23
1	<i>99.94</i>	99.85	99.78	99.73	99.73	99.73	99.73	99.73
2	87.82	<i>99.87</i>	99.84	99.82	99.82	99.82	99.82	99.82
3	85.20	97.27	<i>99.72</i>	99.71	99.71	99.70	99.70	99.70
5	82.59	95.39	98.17	<i>99.68</i>	99.68	99.68	99.68	99.68
10	79.42	94.77	96.35	98.79	<i>99.60</i>	99.60	99.60	99.60
15	80.99	94.75	96.44	98.81	99.52	<i>99.64</i>	99.64	99.64
20	80.93	94.79	96.37	98.75	99.51	99.62	<i>99.63</i>	99.63
23	79.88	94.60	96.16	98.75	99.49	99.60	99.62	<i>99.61</i>

errors — but the fact that they are close to 100 indicates that the numerical impact of these additional sources of error is small.

With overhedging due to excess factors, naïve factor hedging performs as if there were no factor cardinality error (see the values to the right of the diagonal in Table 1). This occurs because (i) the physical trading cash flows obtained under the model with DGNF factors are identical to those obtained under the model with K factors, as the rolling intrinsic policy generates the same actions in both cases; (ii) ignoring possible errors in the estimated deltas, overhedging simply involves hedging additional hypothetical (zero probability) futures curve changes, but this has no impact on the effectiveness of hedges against futures curve changes that can occur; and (iii) the differences in the computed deltas turn out to be very small (this is discussed below).

More surprisingly, with the notable exceptions of one, fifteen, twenty, and twenty-three factors used for delta estimation, naïve factor hedging performs disastrously with underhedging due to omitted factors. These HEM estimates are large and negative; i.e., hedging dramatically *increases* the residual cash flow variance relative to the unhedged cash flow variance. It is striking how large the effect on naïve factor hedging is, since the omitted factors only contribute a few percentage points to the futures price variances. Compare the entries along the diagonal in Table 1 to those just below the diagonal when K varies from two to ten: The fraction of the variance explained by adding one factor to the data generating model is just between 1% and 3.5%. In contrast, the HEM estimates drop from the 96.18-99.54 range to very large negative values between $-5,504.48$ and -70.11 .

The disastrous performance of naïve factor hedging with two to ten factors (K) and underhedging due to omitted factors is a consequence of this strategy taking extremely large trading

positions that magnify small factor cardinality errors. Table 2 displays the average cumulative long and short futures replicating positions taken by naïve factor hedging along a price path for the combinations of DGNF and K values in Table 1. The positions with K between two and ten are large compared to the positions with K less than two or more than ten. These large position sizes are induced by the small determinants of the position generating matrices in the denominator of equation (11). Without factor cardinality error, these positions, even if very large, replicate well the value of storage because they are consistent with price changes from the data generating model (although they might be impractical to trade). With underhedging due to omitted factors, however, these positions are structurally inconsistent with the realized price changes, and their magnitudes magnify the omitted randomness in the realized price changes. This generates severe discrepancies between the realized physical trading cash flows and the replicated cash flows.

The exceptions to the poor performance in Table 1 also deserve comment. When K is equal to one, the stated matrix inversion reduces to a division by a scalar and does not generate huge trading positions. When K is equal to fifteen, the trading positions can be large in size but the empirically calibrated model error is extremely small. When K is equal to twenty or twenty-three, the stated matrix inversion does not appear to be problematic, as factor hedging converges to bucket hedging when the value of K approaches the number of maturities in the term of the storage contract and, as discussed below, bucket hedging is remarkably robust to model error.

Fine-tuned factor hedging. The poor performance of naïve factor hedging with underhedging due to omitted factors motivated us to develop our fine-tuned factor hedging strategy. Table 3 displays the HEM estimates for fine-tuned factor hedging. With no factor cardinality error or with excess factors, fine-tuned factor hedging performs extremely well, as expected, with HEM estimates above 99.60. Although we expect fine-tuned factor hedging to perform better than naïve factor hedging with underhedging due to omitted factors, how much is unclear a priori.

Our results indicate that the improvement is dramatic. Specifically, the HEM estimates are 94.60 or higher with K values equal to two or greater. Fine-tuned factor hedging weakly dominates naïve factor hedging in all cases with factor cardinality error due to omitted factors. In fact, it achieves near optimal performance. Moreover, for a given DGNF value, the relevant HEM estimates increase in the K values, becoming essentially equal to 100 when K meets or exceeds DGNF. The increase is due to the obvious reduction in factor cardinality error. For a given value of K , the HEM estimates generally decrease when the value of DGNF increases

Table 4: HEM estimates for naïve factor hedging, fine-tuned factor hedging, and bucket hedging with model class error.

K							
1	2	3	5	10	15	20	23
Naïve Factor hedging							
89.81	-18784.85	-2591.03	-9540.92	-946.06	-210.13	99.81	99.79
Fine-Tuned Factor Hedging							
89.80	94.91	98.79	98.96	99.60	99.67	99.79	99.80
Bucket Hedging							
99.65	99.80	99.81	99.81	99.80	99.80	99.80	99.80

and factor cardinality error caused by omitted factors worsens.

The improvement in performance of fine-tuned factor hedging relative to naïve factor hedging is due to the smaller size of the replicating positions of fine-tuned factor hedging: The average cumulative long trading positions vary from 10.60 to 28.99 across the considered pairs of DGNF and K values; the range for the average cumulative short trading positions is 0.66 to 17.85 (we omit the table for brevity). This underscores the impact of carefully selecting the position generation matrix when implementing factor hedging.

Bucket hedging. The HEM estimates of bucket hedging vary between 99.60 and 99.92 across the considered pairs of DGNF and K values (we omit the table for brevity). Consistent with the bounds established in Proposition 3, the smallest and largest realized positions of bucket hedging (i.e., deltas) across all sample paths, trading dates, and contracts traded are -0.30 and 0.59 irrespective of the DGNF and K values. The average cumulative long and short trading positions vary in between 25.86 and 26.36 and 14.31 and 15.18, respectively. With factor cardinality error due to omitted factors, bucket hedging outperforms fine-tuned factor hedging, with all the HEM estimates above 99.60, irrespective of the number of factors used to generate the data. Without such error or with factor cardinality error due to excess factors, the performances of bucket hedging and fine-tuned factor hedging are equivalent.

The very similar performance of bucket hedging across different cases is due to the deltas computed in each of these cases being essentially equal to those computed without factor cardinality error. This is remarkable, especially as bucket hedging requires no optimization. Indeed, in the presence of factor cardinality error due to omitted factors, bucket hedging can be intuitively interpreted as the “maximally robust” delta hedging method conditional on the deltas being exactly estimated, as bucket hedging makes no assumptions on the value of DGNF.

The difference in performance between naïve factor hedging and bucket hedging can be explained by the fact that model error can affect hedging performance through two channels:

First, model error can lead to incorrectly computed option deltas. However, this effect is present in both bucket hedging and factor hedging, and our delta estimates appear accurate, as suggested by the near optimal performance of bucket hedging. Second, factor hedging attempts to hedge common randomness across futures prices by relying on specific assumptions about the correlation structure of that randomness. If actual futures price dynamics include additional or different types of randomness which the assumed model structure omits, this can cause hedging errors. Moreover, extreme contract positions tend to magnify the impact of unmodeled randomness in futures prices. The positions in naïve factor hedging can be extremely large, due to the very small determinants of its chosen position generating matrices; hence, even statistically small amounts of omitted randomness can be magnified to produce large hedging errors. In contrast, given the option deltas, bucket hedging does not impose any additional specific structure on the futures curve dynamics; instead, each dated cash flow is hedged individually using futures contracts with the matching maturity. Consequently, bucket hedging is an effective approach to reduce the impact of model error on hedging because this second channel of price correlation error is absent in bucket hedging.

Model Class Error. Similar to our valuation analysis, the particular correlation structure of the factor models used to generate the simulated data may affect our factor cardinality error hedging results. Hence, we conduct a model class error analysis analogous to the one in §6.3. That is, we generate futures curve paths by simulating our calibrated price model with four factors and maturity-specific idiosyncratic noise. This data generating model retains the variance structure of the full twenty-three factor model but has a dissimilar correlation structure.

Table 4 reports the HEM estimates for our three hedging methods. The results in Table 4 are generally consistent with the ones obtained in this analysis with the value of DGNF equal to twenty-three. One difference is that the HEM estimates in Table 4 do not always improve with increasing K values. Compare, for instance, the HEM estimates with naïve factor hedging when the value of K increases from two to five. We explain this observed behavior as follows: As K increases the model used for valuation and hedging captures more noise, but the hedging method includes more replicating positions exposed to noise, so the net effect is unclear.

Overall, these results suggest that our insights obtained from our factor cardinality error analysis do not appear to depend in a critical manner on the particular correlation structure of the omitted randomness.

Table 5: BHEM values for the three hedging approaches.

K	In-Sample			Out-of-Sample		
	Naïve	Fine-Tuned	Bucket	Naïve	Fine-Tuned	Bucket
	Factor	Factor		Factor	Factor	
Hedging	Hedging	Hedging	Hedging	Hedging	Hedging	
1	63.30	60.89	99.59	56.04	51.89	99.81
2	-4283.75	96.09	99.57	-139.33	90.93	99.81
3	-3131.44	99.33	99.65	-1649.49	98.26	99.77
5	-1890.47	99.58	99.72	-4985.42	99.72	99.75
10	-385.49	99.77	99.73	-194.62	99.76	99.74
15	-78.92	99.75	99.74	-2063.01	99.72	99.74
20	96.00	99.76	99.74	83.30	99.71	99.74
23	99.74	99.74	99.74	99.74	99.74	99.74

6.6 Hedging Results with Real Data

We conclude with backtested hedging results obtained on real historical futures data. Since the true price model is unknown, our backtesting analysis reflects some combination of model class error, factor cardinality error, and parameter error. Because history contains a single “sample path,” we split the data into rolling two-year subperiods to obtain multiple subsamples. We perform nine in-sample experiments (for subperiods 1997-1998, 1998-1999, . . . , 2005-2006) and seven out-of-sample experiments (for 1999-2000 through 2005-2006 where the model is estimated over the preceding two years).

We expect the performance of the various hedging approaches to be comparable or somewhat inferior, due to the possible presence of parameter error, to their performance with simulated data. Indeed, our results, reported in Table 5, largely confirm those obtained with simulated data: Naïve factor hedging can perform disastrously, fine-tuned factor hedging is effective once the most important statistical factors are included, and bucket hedging is yet more effective. Thus, the latter two approaches are effective at mitigating the negative effect of various types of model errors on hedging. Moreover, the possible presence of parameter error does not change this conclusion.

6.7 Summary, Managerial Principles, and Practical Implications

Our numerical results show that futures price models with the most important three to six statistical factors are effective at valuing and hedging twenty-three month natural gas storage contracts, yielding valuation and hedging performances comparable to those obtained with full dimensional price models, provided that fine-tuned factor hedging or bucket hedging are

employed. Naïve factor hedging is extremely sensitive to statistically small model errors, while our fine-tuned factor hedging method is rather insensitive to such errors and is near optimal, provided that the most important statistical factors are included in the valuation and hedging model. The performance of bucket hedging is remarkably robust to factor cardinality and model class error and is near optimal.

Our results suggest some general managerial principles on the use of multifactor models to value and hedge other commodity conversion assets and financial options that depend on term structures, including commodity shipping and refining assets, commodity and energy swing and put-call Bermudan options, mortgages, and interest rate caps and floors, and options on mortgages and interest rates term structures in the presence of model error. (i) For valuation, incorporating the most important statistical factors in the term structure model appears satisfactory. (ii) For hedging, (iia) bucket hedging, which ignores the factor correlation structure, is very robust; (iib) but factor hedging, which, by considering the factor correlation structure, trades fewer contracts and may be more convenient than bucket hedging, should be implemented by choosing those futures contracts that yield minimum size positions.

Our valuation results contrast with the apparent widespread use among commodity storage traders of full dimensional price models (Lai et al. 2010). These traders likely prefer the simpler bucket hedging to factor hedging. However, bucket hedging and the number of factors in the valuation model are separable, i.e., all of the bucket hedges use all the relevant futures contracts, irrespective of the number of factors in the underlying valuation model. Our experiments show that bucket hedging based on a lower dimensional price model gives similar hedging results. This suggests there may be redundancy in current commodity storage practice in the number of factors used for valuation, but not in the number of contracts used for hedging. Fine-tuned factor hedging may offer an appealing alternative to bucket hedging when trading in all the contracts in the term of a long dated (storage real) option is not desirable or viable, e.g., due to reduced liquidity.

7. Conclusions

Model error is an important concern in practice when merchants value and hedge the cash flows of commodity and energy conversion assets, as well as financial options. In this paper we study the impact of various types of empirically calibrated term structure model errors on the merchant management of assets that store natural gas, an important energy source. Our research brings to light the differential effect of these errors on the valuation and hedging of

these assets, and proposes a near optimal approach to mitigate their negative effect on the performance of factor hedging. Besides natural gas storage, our research suggests managerial principles to deal with model error when using multifactor models to value and hedge other commodity conversion assets and financial options that are contingent on futures term structure dynamics. For future work, it would be interesting to extend our delta hedging analysis to other hedging approaches, e.g., gamma and vega hedging, and to study the effect of low liquidity on different hedging methods when hedging long-dated real options.

Acknowledgments

François Margot was supported by ONR grant N00014-09-1-0033 and NSF grant CMMI1129163.

Nicola Secomandi was supported by NSF grant CMMI1129163.

References

- Birge, J. R. 2000. Option methods for incorporating risk into linear capacity planning models. *Manufacturing & Service Operations Management* **2**(1) 19–31.
- Bjerkstrand, P., G. Stensland, F. Vagstad. 2011. Gas storage valuation: Price modeling v. optimization models. *The Energy Journal* **32**(1) 203–227.
- Black, F. 1976. The pricing of commodity contracts. *Journal of Financial Economics* **3**(1-2) 167–179.
- Blanco, C., D. Soronow, P. Stefiszyn. 2002. Multi-factor models for forward curve analysis: An introduction to Principal Component Analysis. *Commodities Now* June 76–78.
- Boogert, A., C. de Jong. 2008. Gas storage valuation using a Monte Carlo method. *Journal of Derivatives* **15**(3) 81–98.
- Boogert, A., C. de Jong. 2011/12. Gas storage valuation using a multifactor price process. *Journal of Energy Markets* **4**(4) 29–52.
- Borovkova, S., H. Geman. 2006. Seasonal and stochastic effects in commodity forward curves. *Review of Derivatives Research* **9**(2) 167–186.
- Borovkova, S., H. Geman. 2008. Forward curve modeling in commodity markets. H. Geman, ed., *Risk Management in Commodity Markets: From Shipping to Agricultural and Energy*. John Wiley & Sons, Chichester, UK, 9–32.
- Boyabatli, O., P. R. Kleindorfer, S. R. Koontz. 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science* **57**(10) 1771–1787.
- Brace, A., D. Gatarek, M. Musiela. 1997. The market model of interest rate dynamics. *Mathematical Finance* **7**(2) 127–147.
- Broadie, M., P. Glasserman. 1996. Estimating security price derivatives using simulation. *Management Science* **42**(2) 269–285.
- Caldentey, R., M. Haugh. 2006. Optimal control and hedging of operations in the presence of financial markets. *Mathematics of Operations Research* **31**(2) 285–304.
- Carmona, R., M. Ludkovski. 2010. Valuation of energy storage: An optimal switching approach. *Quantitative Finance* **10**(4) 359–374.
- Casassus, J., P. Collin-Dufresne. 2005. Stochastic convenience yields implied from interest rates and commodity futures. *Journal of Finance* **60**(5) 2283–2331.

- 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.
- Clelow, L., C. Strickland. 2000. *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, London, UK.
- Cortazar, G., E. S. Schwartz. 1994. The valuation of commodity-contingent claims. *Journal of Derivatives* **1**(4) 27–39.
- Detemple, J. 2006. *American-style Derivatives: Valuation and Computation*. Taylor & Francis, London, UK.
- Devalkar, S. K., R. Anupindi, A. Sinha. 2011. Integrated optimization of procurement, processing and trade of commodities in a network environment. *Operations Research* **59**(6) 1369–1381.
- Ding, Q., L. Dong, P. Kouvelis. 2007. On the integration of production and financial hedging decisions in global markets. *Operations Research* **55**(3) 470–489.
- Driessen, J., P. Klaassen, B. Melenberg. 2003. The performance of multi-factor term structure models for pricing and hedging caps and swaptions. *Journal of Financial and Quantitative Analysis* **38**(3) 635–678.
- EIA. 2010. Annual Energy Review 2009. Report #: DOE/EIA-0384(2009), U.S. Energy Information Administration, Washington, DC, USA.
- EIA. 2011a. Annual Energy Outlook 2011 with Projections to 2035. Report #: DOE/EIA-0383(2011), U.S. Energy Information Administration, Washington, DC, USA.
- EIA. 2011b. Natural gas consumption by end use. http://205.254.135.24/dnav/ng/ng_cons_sum_dcu_nus_a.htm.
- EIA. 2011c. Peak underground working natural gas storage capacity. <http://www.eia.gov/naturalgas/storagecapacity/>.
- Eydeland, A., K. Wolyniec. 2003. *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. John Wiley & Sons, Inc., Hoboken, NJ, USA.
- Fan, R., A. Gupta, P. Ritchken. 2001. On pricing and hedging in the swaption market: How many factors, really? Working Paper, Case Western Reserve University, Cleveland, OH, USA.
- Fu, M. C., J. Hu. 1995. Sensitivity analysis for Monte Carlo simulation of option pricing. *Probability in the Engineering and Information Sciences* **9**(3) 417–449.
- Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing & Service Operations Management* **7**(2) 103–120.
- Geman, H. 2005. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, Chichester, UK.
- Geman, H., V.-N. Nguyen. 2005. Soybean inventory and forward curve dynamics. *Management Science* **51**(7) 1076–1091.
- Gibson, R., E. S. Schwartz. 1990. Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance* **45**(3) 959–976.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. Springer, New York, NY, USA.
- Graves, S., H. Meal, S. Dasu, Y. Qiu. 1986. Two-stage production planning in a dynamic environment. S. Axsäter, C. Schneeweiss, E. Silver, eds., *Multi-Stage Production Planning and Inventory Control*. Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin, Germany, 9–43.
- Gray, J., P. Khandelwal. 2004a. Realistic gas storage models II: Trading strategies. *Commodities Now* September 1–5.
- Gray, J., P. Khandelwal. 2004b. Towards a realistic gas storage model. *Commodities Now* June 1–4.
- Heath, D., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions* **26**(3) 17–30.
- Hull, J. C. 2010. *Risk Management and Financial Institutions*. 2nd ed. Prentice Hall, Boston, MA, USA.

- Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50**(7) 909–921.
- Kennedy, D. P. 1994. The term structure of interest rates as a Gaussian random field. *Mathematical Finance* **4**(3) 247–258.
- 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.
- Leppard, S. 2005. *Energy Risk Management: A Non-Technical Introduction to Energy Derivatives*. Risk Books, London, UK.
- Longstaff, F. A., P. Santa-Clara, E. S. Schwartz. 2001. The relative valuation of caps and swaptions: Theory and empirical evidence. *Journal of Finance* **56**(6) 2067–2109.
- Maragos, S. 2002. Valuation of the operational flexibility of natural gas storage reservoirs. E. Ronn, ed., *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 431–456.
- Pilipovic, D. 2007. *Energy Risk: Valuing and Managing Energy Derivatives*. 2nd ed. McGraw-Hill, New York, NY, USA.
- Ross, S. A. 1997. Hedging long run commitments: Exercises in incomplete market pricing. *Economic Notes* **26**(2) 385–420.
- Routledge, B. R., D. J. Seppi, C. S. Spatt. 2000. Equilibrium forward curves for commodities. *Journal of Finance* **55**(3) 1297–1338.
- Schwartz, E. S. 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance* **52**(3) 923–973.
- Schwartz, E. S., J. E. Smith. 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46**(7) 893–911.
- Secomandi, N. 2010a. On the pricing of natural gas pipeline capacity. *Manufacturing and Service Operations Management* **12**(3) 393–408.
- Secomandi, N. 2010b. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.
- Secomandi, N., M. X. Wang. 2012. A computational approach to the real option management of network contracts for natural gas pipeline transport capacity. *Manufacturing & Service Operations Management* **14**(3) 441–454.
- Seppi, D. J. 2002. Risk-neutral stochastic processes for commodity derivative pricing: An introduction and survey. E. Ronn, ed., *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 3–60.
- Smith, J. E., K. F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research* **47**(1) 1–15.
- Suenaga, H., A. Smith, J. Williams. 2008. Volatility dynamics of NYMEX natural gas futures prices. *Journal of Futures Markets* **28**(5) 438–463.
- Thompson, M. 2012. Natural gas storage valuation, optimization, market and credit risk management. Working Paper, Queen’s School of Business, Queen’s University, Kingston, Ontario, CA.
- Thompson, M., M. Davison, H. Rasmussen. 2009. Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics* **56**(3) 226–238.
- Tirole, J. 2006. *The Theory of Corporate Finance*. Princeton University Press, Princeton, NJ, USA.
- Tolmasky, C., D. Hindanov. 2002. Principal components analysis for correlated curves and seasonal commodities: The case of the petroleum market. *Journal of Futures Markets* **2**(11) 1019–1035.
- Van Mieghem, J. A. 2003. Capacity management, investment and hedging: Review and recent developments. *Manufacturing & Service Operations Management* **5**(4) 269–302.
- Veronesi, P. 2010. *Fixed Income Securities: Valuation, Risk, and Risk Management*. John Wiley & Sons Inc., Hoboken, NJ, USA.
- Wu, O., H. Chen. 2010. Optimal control and equilibrium behavior of production-inventory systems. *Management Science* **56**(8) 1362–1379.

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

Online Appendix

A. Illustration of Merchant Commodity Storage

A simple two-date binomial example illustrates the two parts of merchant commodity storage:

(i) Risk neutral valuation of the storage cash flows and (ii) delta hedging.

For concreteness, assume that one unit of a commodity can be stored from date 0 to date 1. Denote a feasible storage policy by π . Assuming no inventory in storage at date 0, this policy specifies (i) the amount of commodity purchased at date 0 and stored until date 1, given the date 0 spot price, s_0 , and the date 0 futures price for date 1 delivery, $F_{0,1}$, and (ii) any amount of commodity sold at date 1 given the date 1 spot price, s_1 . Denote by CF_0^π the date 0 cash flow of this policy and by \widetilde{CF}_1^π the date 1 *random* cash flow of this policy; \widetilde{CF}_1^π is random because it depends on the date 1 random spot price of the commodity, \tilde{s}_1 .

The market value of the random date 1 cash flow of this policy, \widetilde{CF}_1^π , can be determined by replication, which is mathematically equivalent to an approach called *risk neutral valuation* (Luenberger 1998, pages 327-330, Smith and McCardle 1999). The idea is to set up at date 0 a portfolio consisting of B^π dollars worth of a risk-less bond and a position q^π in futures contracts with date 1 delivery so that the date 1 cash flow of this portfolio exactly replicates the date 1 cash flow of the policy π for each realization of the date 1 spot price. The date 0 market value of this *replicating portfolio* is thus the date 0 market value of the random date 1 cash flow of this policy, \widetilde{CF}_1^π . Specifically, the market value of this portfolio is B^π because the value (rather than the price) of a futures position by definition is zero at inception at date 0.

For simplicity, assume that the risk-free interest rate is zero. The date 1 cash flow from the bond position B^π taken at date 0 is thus just B^π . The date 1 random cash flow from taking the position q^π in futures contracts with date 1 delivery is $(\tilde{s}_1 - F_{0,1})q^\pi$. Again for simplicity, consider the following binomial commodity price model (Hull 2012, p. 368): Let u be a number greater than 1 and suppose that the date 1 spot price can either be equal to $uF_{0,1}$ or $F_{0,1}/u$. Denote by $CF_1^{\pi,u}$ and $CF_1^{\pi,1/u}$ the date 1 cash flows of the policy π when the date 1 spot price is equal to $uF_{0,1}$ and $F_{0,1}/u$, respectively.

The portfolio replicating conditions are

$$(uF_{0,1} - F_{0,1})q^\pi + B^\pi = CF_1^{\pi,u}, \quad (13)$$

$$(F_{0,1}/u - F_{0,1})q^\pi + B^\pi = CF_1^{\pi,1/u}. \quad (14)$$

The solution to these equations is

$$q^\pi = \frac{CF_1^{\pi,u} - CF_1^{\pi,1/u}}{uF_{0,1} - F_{0,1}/u}, \quad (15)$$

$$B^\pi = \frac{1}{1+u}CF_1^{\pi,u} + \frac{u}{1+u}CF_1^{\pi,1/u}. \quad (16)$$

Since $u > 1$ and $1/(1+u) + u/(1+u) \equiv 1$, the terms $1/(1+u)$ and $u/(1+u)$ in (16) can be interpreted as probabilities. These are called the risk neutral probabilities of the events that the date 1 random spot price is equal to $uF_{0,1}$ and $F_{0,1}/u$, respectively (Hull 2012, p. 368). We denote by \mathbb{E} expectation taken using these risk neutral probabilities.

We can thus express the date 0 market value B^π of the random cash flow \widetilde{CF}_1^π as the date 0 risk neutral expected value (no discounting is applied because the risk-free rate is assumed to be zero):

$$B^\pi \equiv \mathbb{E} \left[\widetilde{CF}_1^\pi \right]. \quad (17)$$

This is the risk neutral valuation approach. Although explained here in the context of valuing the random storage cash flow \widetilde{CF}_1^π , risk neutral valuation applies to the valuation of any other date 1 random cash flows whose randomness is solely due to their dependence on the date 1 random spot price.

The date 0 market value of the storage policy π , V_0^π , is the sum of the date 0 cash flow and the risk neutral expected value of the date 1 cash flow of this policy:

$$V_0^\pi = CF_0^\pi + \mathbb{E} \left[\widetilde{CF}_1^\pi \right].$$

An optimal storage policy maximizes the market value of the storage capacity. That is, it solves the optimization problem

$$\max_{\pi} V_0^\pi. \quad (18)$$

Futures prices are risk neutral expected future spot prices (Hull 2012, p. 369):

$$\mathbb{E} [\tilde{s}_1] = \frac{1}{1+u}uF_{0,1} + \frac{u}{1+u}\frac{F_{0,1}}{u} = F_{0,1}.$$

Assume that $F_{0,1} > s_0$ and that commodity trading and storage incur no commodity loss and inventory adjustment cost. The optimal policy thus purchases one unit of commodity at date 0, stores it until date 1, and sells it at date 1, and its date 0 and date 1 cash flows are $CF_0^* := -s_0$ and $\widetilde{CF}_1^* := \tilde{s}_1$, respectively. The date 0 market value of these cash flows, and hence of this policy, is

$$V_0 := CF_0^* + \mathbb{E} \left[\widetilde{CF}_1^* \right] = -s_0 + \mathbb{E} [\tilde{s}_1] = F_{0,1} - s_0.$$

Suppose that the commodity storage capacity can be bought at the discounted price of αV_0 , where $\alpha \in (0, 1)$. Consider a merchant that purchases this capacity for αV_0 and then follows the optimal storage policy that produces cash flows with an ex-ante market value V_0 . The *positive* market value of this trade is $(1 - \alpha)V_0$.

Although this strategy has positive value, the merchant's realized profit (or loss) with this strategy is random. In particular, the merchant's date 1 possible trading loss is

$$-\min \left\{ \text{CF}_0^* + \widetilde{\text{CF}}_1^*, 0 \right\},$$

which we express as a nonnegative random variable. We assume that the merchant, when implementing this risky strategy, incurs at date 0 *risk capital charges proportional*, according to $\eta > 0$, to the date 0 market value of any date 1 trading losses. The merchant can use financial hedging to create an additional cash flow at date 1 to reduce, ideally eliminate, these risk capital charges. Specifically, at date 0 the merchant can short a number q of futures contracts with date 1 delivery. In this case the merchant's risk capital charges are

$$\text{RCC}_0^q := \eta \mathbb{E} \left[-\min \left\{ \text{CF}_0^* + \widetilde{\text{CF}}_1^* - (\tilde{s}_1 - F_{0,1})q, 0 \right\} \right].$$

The replicating conditions (13)-(14) and the condition (17) applied to the optimal policy, that is, $q^* := (\text{CF}_1^{*,u} - \text{CF}_1^{*,1/u}) / [F_{0,1}(u - 1/u)]$ and $B^* := \mathbb{E} [\widetilde{\text{CF}}_1^*]$, imply that the date 1 payoff from taking a position equal to q^* futures contracts at date 0, that is, $(\tilde{s}_1 - F_{0,1})q^*$, satisfies

$$(\tilde{s}_1 - F_{0,1})q^* = \widetilde{\text{CF}}_1^* - B^* = \widetilde{\text{CF}}_1^* - \mathbb{E} [\widetilde{\text{CF}}_1^*].$$

Therefore, by *shorting* q^* futures contracts at date 0, the merchant's date 1 total payoff, which includes the date 1 optimal storage cash flow, $\widetilde{\text{CF}}_1^*$, and the cash flow from financial hedging, $-(\tilde{s}_1 - F_{0,1})q^*$, is the *non risky* (non random) quantity

$$\widetilde{\text{CF}}_1^* - (\tilde{s}_1 - F_{0,1})q^* = \widetilde{\text{CF}}_1^* - \widetilde{\text{CF}}_1^* + \mathbb{E} [\widetilde{\text{CF}}_1^*] = \mathbb{E} [\widetilde{\text{CF}}_1^*].$$

Recalling that V_0 is positive, it follows that

$$-\min \left\{ \text{CF}_0^* + \widetilde{\text{CF}}_1^* - (\tilde{s}_1 - F_{0,1})q^*, 0 \right\} = -\min \left\{ \text{CF}_0^* + \mathbb{E} [\widetilde{\text{CF}}_1^*], 0 \right\} = -\min \{V_0, 0\} = 0.$$

This financial hedging strategy thus eliminates the merchant's date 0 risk capital charges, because there is no possibility of a date 1 negative payoff. That is, the merchant is, by construction, perfectly hedged and the realized payoff of the combined storage and financial hedging strategy is $(1 - \alpha)V_0$ *irrespective* of the realization of the date 1 random spot price.

This particular financial hedging strategy is known as *delta hedging* because the shorted futures position q^* is called the time 0 delta of the optimal value of the storage capacity, denoted by Δ_0 . Since delta hedging makes the risk capital charges equal to 0 *with probability 1*, it optimally solves the problem of minimizing the market value of the risk capital charges when the optimal storage policy is used:

$$\min_q \text{RCC}_0^q. \quad (19)$$

In this sense, delta hedging is an *optimal* financial hedging policy.

We have thus shown that jointly determining an optimal storage policy and an optimal financial hedging policy reduces to solving two simpler and related problems: (i) Determine an optimal storage policy by solving problem (18); (ii) compute the delta of this policy, Δ_0 , and short an equal number of futures contracts, thus optimally solving problem (19). That is, this two-step approach optimally solves the following problem, where $\text{RCC}_0^{\pi,q}$ are the risk capital charges incurred when following the storage policy π and the financial hedging policy q :

$$\max_{\pi,q} (1 - \alpha)V_0^\pi - \text{RCC}_0^{\pi,q}.$$

An analogous approach remains valid when determining an optimal storage policy is impractical, as is the case in the setting of this paper. That is, suppose that the merchant can only compute a *feasible*, rather than optimal, storage policy π with *positive* market value V_0^π and can purchase the storage capacity at the discounted value αV_0^π . The merchant can eliminate the risk capital charges $\text{RCC}_0^{\pi,q}$ by shorting at date 0 a number of futures contracts equal to q^π as given by (15), which is called the delta of the policy π and is denoted by Δ_0^π . In other words, shorting the delta Δ_0^π is an optimal financial hedging policy *conditional* on the storage policy π being used, in the sense that it optimally solves the following problem:

$$\min_q \text{RCC}_0^{\pi,q}.$$

When purchasing the storage capacity for αV_0^π and managing storage according to the policy π , shorting a number of futures contracts equal to the delta Δ_0^π ensures that the merchant earns the positive profit $(1 - \alpha)V_0^\pi$ with probability 1.

Delta hedging is also a *practical* policy. The delta Δ_0^π can be easily computed when using the simple binomial model. As discussed in this paper, under more advanced continuous time and state models of commodity price evolution, the date 0 delta of a feasible storage policy is equal to the partial derivative of the value of this policy with respect to the futures price $F_{0,1}$.

When no closed-form expressions exist for this derivative, delta can be estimated numerically. Moreover, in theory delta, and hence the futures hedging position, must be recomputed and rebalanced continuously over time for delta hedging to eliminate the risk capital charges fully. In practice, however, periodic rebalancing is sufficient for delta hedging to be effective. Delta hedging also applies when storage involves random cash flows at more than one date, in which case there is a separate delta for each future date with a random cash flow. Finally, delta hedging applies regardless of the particular commodity conversion asset being managed. It is thus relevant for general merchant operations problems.

B. Deltas with Exact Storage Valuation

Proposition 6 in this appendix extends the pathwise delta characterization established in Proposition 2 to the case of exact storage valuation based on the following MDP for exact storage valuation (Lai et al. 2010, §2):

$$V_n(x_n, \mathbf{F}_n) = \max_{a \in \mathcal{A}(x_n)} p(a, s_n) + \delta \mathbb{E} [V_{n+1}(x_n - a, \mathbf{F}_{n+1}) | \mathbf{F}'_n], \quad (20)$$

for all $n \in \mathcal{N}$ and $(x_n, \mathbf{F}_n) \in \mathcal{X} \times \mathfrak{R}_+^{N-n}$, where $V_n(x_n, \mathbf{F}_n)$ denotes the optimal value function in stage n and state (x_n, \mathbf{F}_n) , and $V_N(x_N, \mathbf{F}_N) := 0$, for all $x_N \in \mathcal{X}$ (recall that $\mathbf{F}_N \equiv 0$).

Proposition 4 characterizes the optimal value function and an optimal operating policy of this MDP. The proof of this result is a simple adaptation of the proofs of Lai et al. (2010, Theorem 1) and Secomandi (2010, Theorem 1).

Proposition 4 (Concavity and basestock optimality). *In every stage n , the function $V_n(x_n, \mathbf{F}_n)$ is concave in x_n for each given \mathbf{F}_n , and the optimal policy for the MDP (20) features two basestock targets, $\underline{b}_n(\mathbf{F}_n), \bar{b}_n(\mathbf{F}_n) \in \mathcal{X}$, such that $\underline{b}_n(\mathbf{F}_n) \leq \bar{b}_n(\mathbf{F}_n)$ and an optimal decision rule $A_n^*(x_n, \mathbf{F}_n)$ satisfies*

$$A_n^*(x_n, \mathbf{F}_n) = \begin{cases} C^I \vee [x_n - \underline{b}_n(\mathbf{F}_n)], & x_n \in [0, \underline{b}_n(\mathbf{F}_n)), \\ 0, & x_n \in [\underline{b}_n(\mathbf{F}_n), \bar{b}_n(\mathbf{F}_n)], \\ C^W \wedge [x_n - \bar{b}_n(\mathbf{F}_n)], & x_n \in (\bar{b}_n(\mathbf{F}_n), \bar{x}]. \end{cases} \quad (21)$$

Proposition 5, based on Lemma 1, establishes basic properties of the optimal value function and basestock targets, which are used to establish Proposition 6. All of these results are based on Assumption 2.

Assumption 2 (Capacities and maximum space). The capacity limits C^I and C^W as well as the maximal inventory level \bar{x} are integer multiples of some positive real number.

We denote by Q the largest common factor of C^I , C^W , and \bar{x} . Lemma 1 is related to Secomandi (2010, Propositions 2 and 3).

Lemma 1 (Characterization). *Under Assumption 2, in every stage $n \in \mathcal{N}$: (a) The function $V_n(x, \mathbf{F}_n)$ is piecewise linear continuous in inventory $x \in \mathcal{X}$ with break points in set $\mathcal{Q} := \{0, Q, 2Q, \dots, \bar{x}\}$ for each given futures curve $\mathbf{F}_n \in \mathfrak{R}_+^{N-n}$; (b) There exist optimal basestock levels $\underline{b}_n(\mathbf{F}_n)$ and $\bar{b}_n(\mathbf{F}_n)$ in set \mathcal{Q} , $\forall \mathbf{F}_n \in \mathfrak{R}_+^{N-n}$.*

Following Lemma 1 we define the finite set of feasible actions at inventory level $x \in \mathcal{X}$ as

$$\begin{aligned} \mathcal{A}'(x) := & \{(x - \bar{x}) \vee C^I, \dots, x - (\lfloor x/Q \rfloor + 2)Q, x - (\lfloor x/Q \rfloor + 1)Q\} \cup \{0\} \\ & \cup \{x - \lfloor x/Q \rfloor Q, x - (\lfloor x/Q \rfloor - 1)Q, \dots, x \wedge C^W\}, \end{aligned}$$

with $\{0\}$ removed if it is a duplicate.

Proposition 5 (Representation and Lipschitz continuity). *Under Assumption 2, in every stage $n \in \mathcal{N}$: (a) For all inventory levels $x \in \mathcal{X}$ and futures curves $\mathbf{F}_n \in \mathfrak{R}_+^{N-n}$ it holds that $V_n(x, \mathbf{F}_n) = \max_{a \in \mathcal{A}'(x)} p(a, s_n) + \delta \mathbb{E}[V_{n+1}(x_n - a, \mathbf{F}_{n+1}) | \mathbf{F}'_n]$; (b) For each given inventory $x \in \mathcal{X}$ it holds that the function $V_n(x, \mathbf{F}_n)$ is Lipschitz continuous in the futures curve $\mathbf{F}_n \in \mathfrak{R}_+^{N-n}$; i.e., there exists $L_n(x) \in \mathfrak{R}_+$ such that $|V_n(x, \mathbf{F}_n^2) - V_n(x, \mathbf{F}_n^1)| \leq L_n(x) \sum_{m=n}^N |F_{n,m}^2 - F_{n,m}^1|$, $\forall \mathbf{F}_n^1, \mathbf{F}_n^2 \in \mathfrak{R}_+^{N-n}$.*

In Proposition 6 we indicate by x_m^* the optimal inventory level in stage $m \in \mathcal{N} \setminus \{0\}$ and by $\Delta_{n,m}(t, x_n, \mathbf{G}_n(t))$ the delta under an optimal policy.

Proposition 6 (Pathwise deltas with an optimal policy). *Under Assumption 2, for every $n \in \mathcal{N} \setminus \{0\}$ it holds that*

$$\begin{aligned} \Delta_{n,m}(t, x_n, \mathbf{G}_n(t)) = & \frac{\bar{\delta}(t, T_m)}{F(t, T_m)} \mathbb{E}[(\phi^I \mathbf{1}\{x_m^* \in [0, \underline{b}_m(\mathbf{F}_m)]\}) + \phi^W \mathbf{1}\{x_m^* \in (\bar{b}_m(\mathbf{F}_m), \bar{x}]\}) \\ & s_m A_m^*(x_m^*, \mathbf{F}_m) | x_n, \mathbf{G}_n(t)], \end{aligned} \quad (22)$$

for all $m \in \mathcal{N}_n$, $x_n \in \mathcal{X}$, $t \in [T_{n-1}, T_n)$, and $\mathbf{G}_n(t) \in \mathfrak{R}_+^{N-n}$.

C. Proofs

Additional Notation. Given n and $m \in \mathcal{N}_{n+1}$, $t \in [T_n, T_{n+1})$, and $t' \in (T_n, T_{n+1}]$, with $t < t'$, we define

$$\beta_m(t, t', \mathbf{Y}) := \exp \left[-\frac{1}{2}(t' - t) \sum_{k \in \mathcal{K}} \sigma_{m,k,n}^2 + \sqrt{t' - t} \sum_{k \in \mathcal{K}} \sigma_{m,k,n} Y_k \right]$$

and $\vec{\beta}(t, t', \mathbf{Y}) := (\beta_m(t, t', \mathbf{Y}), m \in \mathcal{N}_{n+1})$. By (2), we can equivalently express $\mathbf{F}'(t')$ given $\mathbf{F}'(t)$ as $\mathbf{F}'(t)\vec{\beta}(t, t', \mathbf{Y})$ and $\mathbf{F}(T_{n+1})$ given $\mathbf{F}'(t)$ as $\mathbf{F}'(t)\vec{\beta}(t, T_{n+1}, \mathbf{Y})$, where the products are in the componentwise sense. Hence, the equalities

$$\begin{aligned} F(t, T_m) &= \mathbb{E}[F(t', T_m)|F(t, T_m)], \\ \mathbb{E}[F(t', T_m)|F(t, T_m)] &= F(t, T_m)\mathbb{E}[\beta_m(t, t', \mathbf{Y})] \end{aligned}$$

imply that

$$\mathbb{E}[\beta_m(t, t', \mathbf{Y})] = 1. \quad (23)$$

Given $n < N - 1$, when $t = T_n$ and $t' = T_{n+1}$, we abbreviate $\beta_m(T_n, T_{n+1}, \mathbf{Y})$ to $\beta_m^{n, n+1}(\mathbf{Y})$ for all $m \in \mathcal{N}_{n+1}$, and $\vec{\beta}(T_n, T_{n+1}, \mathbf{Y})$ to $\vec{\beta}^{n, n+1}(\mathbf{Y})$. We use this notation in the proofs of Lemma 2 and Proposition 5.

We define by $W_n^\pi(x, \mathbf{F}'_n) := \delta\mathbb{E}[V_{n+1}^\pi(x, \mathbf{F}_{n+1})|\mathbf{F}'_n]$ the policy π continuation value in stage $n < N$ given the inventory level x and the futures curve \mathbf{F}'_n . We use this notation in the proof of Proposition 2. This notation allows us to write

$$\begin{aligned} V_n^\pi(x, \mathbf{F}_n) &= \sum_{m=n}^{N-1} \delta^{m-n} \mathbb{E}[p(A_m^\pi(x_m^\pi, \mathbf{F}_m), s_m)|x, \mathbf{F}_n] \\ &= p(A_n^\pi(x, \mathbf{F}_n), s_n) + \sum_{m=n+1}^{N-1} \delta^{m-n} \mathbb{E}[p(A_m^\pi(x_m^\pi, \mathbf{F}_m), s_m)|x - A_n^\pi(x, \mathbf{F}_n), \mathbf{F}'_n] \\ &= p(A_n^\pi(x, \mathbf{F}_n), s_n) + \delta\mathbb{E}[V_{n+1}^\pi(x - A_n^\pi(x, \mathbf{F}_n), \mathbf{F}_{n+1})|\mathbf{F}'_n] \\ &= p(A_n^\pi(x, \mathbf{F}_n), s_n) + W_n^\pi(x - A_n^\pi(x, \mathbf{F}_n), \mathbf{F}'_n). \end{aligned}$$

We define by $W_n(x, \mathbf{F}'_n) := \delta\mathbb{E}[V_{n+1}(x, \mathbf{F}_{n+1})|\mathbf{F}'_n]$ the optimal continuation value in stage $n < N$ given the inventory level x and the futures curve \mathbf{F}'_n . We use this notation in the proofs of Lemma 1 and Propositions 5 and 6. We denote by $v_n(x, a, \mathbf{F}_n)$ the objective function of the maximization on the right hand side of the MDP (20). We use this notation in the proofs of Propositions 5 and 6 and Lemma 3.

Proof of Proposition 1 (Factor hedging positions). Let $d\mathbf{Z}(t)$ be a K -dimensional vector of independent Brownian motions. Denote by $d\mathbf{F}'(t)$ the vector $(dF(t, T_n), \dots, dF(t, T_{N-1}))$, and by $d\mathbf{F}'^B(t)$ and $d\mathbf{F}'^E(t)$, respectively, the vectors of its components corresponding to maturity labels in sets \mathcal{M}_n and \mathcal{L}_n . Write model (1) in vector notation as $d\mathbf{F}'^B(t)/\mathbf{F}'^B(t) = B_{n-1}d\mathbf{Z}(t)$ and $d\mathbf{F}'^E(t)/\mathbf{F}'^E(t) = E_{n-1}d\mathbf{Z}(t)$, where “/” here indicates componentwise division.

Using the assumption that $|B_{n-1}| \neq 0$, it is not hard to show that

$$d\mathbf{F}'^E(t)/\mathbf{F}'^E(t) = E_{n-1}B_{n-1}^{-1}d\mathbf{F}'^B(t)/\mathbf{F}'^B(t),$$

which can be equivalently written as

$$dF(t, T_l) = \sum_{m \in \mathcal{M}_n} \frac{F(t, l)}{F(t, m)} \frac{\gamma_{l, m, n-1}}{|B_{n-1}|} dF(t, T_m), \quad \forall l \in \mathcal{L}_n. \quad (24)$$

Using (24), the second term on the right hand side of (6) can be written as

$$\begin{aligned} \sum_{m \in \mathcal{N}_n} \Delta_{n, m}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_m) &= \sum_{m \in \mathcal{M}_n} \Delta_{n, m}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_m) \\ &\quad + \sum_{l \in \mathcal{L}_n} \Delta_{n, l}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_l) \\ &= \sum_{m \in \mathcal{M}_n} \Delta_{n, m}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_m) \\ &\quad + \sum_{l \in \mathcal{L}_n} \Delta_{n, l}^\pi(t, x, \mathbf{G}_n(t)) \sum_{m \in \mathcal{M}_n} \frac{F(t, l)}{F(t, m)} \frac{\gamma_{l, m, n-1}}{|B_{n-1}|} dF(t, T_m) \\ &= \sum_{m \in \mathcal{M}_n} \Delta_{n, m}^\pi(t, x, \mathbf{G}_n(t)) dF(t, T_m) \\ &\quad + \sum_{m \in \mathcal{M}_n} \sum_{l \in \mathcal{L}_n} \Delta_{n, l}^\pi(t, x, \mathbf{G}_n(t)) \frac{F(t, l)}{F(t, m)} \frac{\gamma_{l, m, n-1}}{|B_{n-1}|} dF(t, T_m) \\ &= \sum_{m \in \mathcal{M}_n} \kappa_{n, m}(t, x, \mathbf{G}_n(t)) dF(t, T_m), \end{aligned}$$

with

$$\kappa_{n, m}(t, x, \mathbf{G}_n(t)) := \Delta_{n, m}^\pi(t, x, \mathbf{G}_n(t)) + \sum_{l \in \mathcal{L}_n} \Delta_{n, l}^\pi(t, x, \mathbf{G}_n(t)) \frac{F(t, l)}{F(t, m)} \frac{\gamma_{l, m, n-1}}{|B_{n-1}|}.$$

Thus, setting $q_{n, m}(t, x, \mathbf{G}_n(t)) = \kappa_{n, m}(t, x, \mathbf{G}_n(t))$, for all $m \in \mathcal{M}_n$, solves the system of linear equations (10). \square

Lemma 2. Lemma 2 is needed in the proof of Proposition 2.

Lemma 2 (Pathwise derivatives). *Under Assumption 1(a), given $x \in \mathcal{X}$, for every $n \in \mathcal{N} \setminus \{N-1\}$ and $m \in \mathcal{N}_{n+1}$ it holds that*

$$\frac{dV_{n+1}^\pi(x, \mathbf{F}_{n+1})}{dF_{n, m}} = \frac{\partial V_{n+1}^\pi(x, \mathbf{F}_{n+1})}{\partial F_{n+1, m}} \frac{F_{n+1, m}}{F_{n, m}}, \quad (25)$$

$$\frac{\partial \mathbb{E}[V_{n+1}^\pi(x, \mathbf{F}_{n+1}) | \mathbf{F}'_n]}{\partial F_{n, m}} = \mathbb{E} \left[\frac{\partial V_{n+1}^\pi(x, \mathbf{F}_{n+1})}{\partial F_{n, m}} | \mathbf{F}'_n \right]. \quad (26)$$

Proof. We interpret $F_{n,m}$ as a parameter and write $\mathbf{F}_{n+1}(F_{n,m})$ to explicitly indicate the dependence of \mathbf{F}_{n+1} on this parameter. In particular, as $F_{n+1,m} = F_{n,m}\beta_m^{n,n+1}(\mathbf{Y})$, $F_{n+1,m}$ depends on $F_{n,m}$ but every other $F_{n+1,l}$, with $l > m$, does not.

We first show that conditions (A1)-(A4) in Appendix A of Broadie and Glasserman (1996) hold with respect to $\mathbb{E}[V_{n+1}^\pi(x, \mathbf{F}_{n+1}(F_{n,m}))|\mathbf{F}'_n]$.

(A1) Under model (2) the quantity $\partial F_{n+1,l}(F_{n,m})/\partial F_{n,m}$ exists with probability 1 because it is equal to $\beta_m^{n,n+1}(\mathbf{Y})$ if $l = m$ and to 0 when $l > m$.

(A2) Fix x and let $\mathcal{D}_{n+1}^{V^\pi}$ denote the set of futures curves \mathbf{F}_{n+1} at which $V_{n+1}^\pi(x, \mathbf{F}_{n+1})$ is differentiable with respect to each element of \mathbf{F}_{n+1} . Assumption 1(a) and Rademacher's theorem imply that $V_{n+1}^\pi(x, \mathbf{F}_{n+1})$ is differentiable almost everywhere on \mathfrak{R}_+^{N-n-1} with respect to each element of \mathbf{F}_{n+1} . Price model (2) implies that $\mathbb{P}(\mathbf{F}_{n+1}(F_{n,m}) \in \mathcal{D}_{n+1}^{V^\pi}|\mathbf{F}'_n) = 1$, where \mathbb{P} denotes probability, for all $F_{n,m} \in \mathfrak{R}_+$.

(A3) This is Assumption 1(a).

(A4) The random variable $F_{n+1,m}(F_{n,m})$ is almost surely Lipschitz continuous with integrable modulus $\beta_m^{n,n+1}(\mathbf{Y})$ because $|F_{n+1,m}(F_{n,m}^2) - F_{n+1,m}(F_{n,m}^1)| = \beta_m^{n,n+1}(\mathbf{Y})|F_{n,m}^2 - F_{n,m}^1|$, for all $F_{n,m}^1, F_{n,m}^2 \in \mathfrak{R}_+$, and $\mathbb{E}[\beta_m^{n,n+1}(\mathbf{Y})] = 1$. Every other random variable $F_{n+1,l}(F_{n,m})$, with $m < l$, is almost surely Lipschitz with integrable modulus 0 because each such random variable does not depend on $F_{n,m}$.

Following Broadie and Glasserman (1996, p. 280), (25) holds because under conditions (A1)-(A2) the pathwise derivative $dV_{n+1}^\pi(x, \mathbf{F}_{n+1})/dF_{n,m}$ satisfies

$$\begin{aligned} \frac{dV_{n+1}^\pi(x, \mathbf{F}_{n+1})}{dF_{n,m}} &= \sum_{l=n+1}^{N-1} \frac{\partial V_{n+1}^\pi(x, \mathbf{F}_{n+1})}{\partial F_{n+1,m}} \frac{\partial F_{n+1,l}}{\partial F_{n,m}} = \frac{\partial V_{n+1}^\pi(x, \mathbf{F}_{n+1})}{\partial F_{n+1,m}} \beta_m^{n,n+1}(\mathbf{Y}) \\ &= \frac{\partial V_{n+1}^\pi(x, \mathbf{F}_{n+1})}{\partial F_{n+1,m}} \frac{F_{n+1,m}}{F_{n,m}}. \end{aligned}$$

Expression (26) follows from Proposition 1 in Broadie and Glasserman (1996). \square

Proof of Proposition 2 (Pathwise deltas). For simplicity of exposition, we prove the claimed result for $t = 0$ and $n = 1$. The proof for the general case follows similar steps. As $\mathbf{F}'_0 \equiv \mathbf{G}_1(T_0)$, in this proof we use the notation \mathbf{F}'_0 in lieu of the notation $\mathbf{G}_1(T_0)$.

Suppose $m = 1$. Formula (5) and Lemma 2 imply that

$$\frac{\Delta_{1,1}^\pi(0, x_1, \mathbf{F}'_0)}{\bar{\delta}(0, T_1)} = \frac{1}{\bar{\delta}(0, T_1)} \frac{\partial U_1^\pi(0, x_1, \mathbf{F}'_0)}{\partial F_{0,1}} = \frac{\partial \mathbb{E}[V_1^\pi(x_1, \mathbf{F}_1)|\mathbf{F}'_0]}{\partial F_{0,1}} = \mathbb{E} \left[\frac{dV_1^\pi(x_1, \mathbf{F}_1)}{dF_{0,1}} \middle| \mathbf{F}'_0 \right]. \quad (27)$$

It follows from Lemma 2 that

$$\frac{dV_1^\pi(x_1, \mathbf{F}_1)}{dF_{0,1}} = \sum_{n=2}^{N-1} \frac{\partial V_1^\pi(x_1, \mathbf{F}_1)}{\partial F_{1,m}} \frac{\partial F_{1,n}(F_{0,1})}{\partial F_{0,1}} = \frac{\partial V_1^\pi(x_1, \mathbf{F}_1)}{\partial F_{1,1}} \beta_1^{0,1}(\mathbf{Y}) = \frac{\partial V_1^\pi(x_1, \mathbf{F}_1)}{\partial F_{1,1}} \frac{s_1}{F_{0,1}}. \quad (28)$$

Pick a vector $\bar{\mathbf{F}}_1$ at which $V_1^\pi(x_1, \mathbf{F}_1)$ is differentiable. It follows from Assumption 1(b) that

$$\left. \frac{\partial V_1^\pi(x_1, \mathbf{F}_1)}{\partial s_1} \right|_{\mathbf{F}_1 = \bar{\mathbf{F}}_1} = \left. \frac{\partial p(a_1^\pi(x_1, \bar{\mathbf{F}}_1), s_1)}{\partial s_1} \right|_{s_1 = \bar{s}_1} + \left. \frac{\partial W_1^\pi(x_1 - a_1^\pi(x_1, \bar{\mathbf{F}}_1), \mathbf{F}'_1)}{\partial s_1} \right|_{\mathbf{F}'_1 = \bar{\mathbf{F}}'_1}. \quad (29)$$

The first term on the right hand side of (29) can be expressed as follows:

$$\begin{aligned} \left. \frac{\partial p(a_1^\pi(x_1, \bar{\mathbf{F}}_1), s_1)}{\partial s_1} \right|_{s_1 = \bar{s}_1} &= \left. \frac{\partial(\phi^I s_1 + c^I) a_1^\pi(x_1, \bar{\mathbf{F}}_1)}{\partial s_1} \right|_{s_1 = \bar{s}_1} \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) < 0\} \\ &\quad + \left. \frac{\partial(\phi^W s_1 - c^W) a_1^\pi(x_1, \bar{\mathbf{F}}_1)}{\partial s_1} \right|_{s_1 = \bar{s}_1} \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) > 0\} \\ &= (\phi^I \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) < 0\} + \phi^W \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) > 0\}) a_1^\pi(x_1, \bar{\mathbf{F}}_1). \end{aligned} \quad (30)$$

As \mathbf{F}'_1 does not depend on $F_{1,1}$, the second term on the right hand side of (29) is zero:

$$\left. \frac{\partial W_1^\pi(x_1 - a_1^\pi(x_1, \bar{\mathbf{F}}_1), \mathbf{F}'_1)}{\partial s_1} \right|_{\mathbf{F}'_1 = \bar{\mathbf{F}}'_1} = 0. \quad (31)$$

It follows from (30) and (31) that (29) can be expressed as

$$\left. \frac{\partial V_1^\pi(x_1, \mathbf{F}_1)}{\partial F_{1,n}} \right|_{\mathbf{F}_1 = \bar{\mathbf{F}}_1} = (\phi^I \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) < 0\} + \phi^W \mathbf{1}\{a_1^\pi(x_1, \bar{\mathbf{F}}_1) > 0\}) a_1^\pi(x_1, \bar{\mathbf{F}}_1). \quad (32)$$

Expressions (27), (28), and (32) imply that

$$\Delta_{1,1}^\pi(0, x_1, \mathbf{F}'_0) = \frac{\bar{\delta}(0, T_1)}{F_{0,1}} \mathbb{E} [(\phi^I \mathbf{1}\{a_1^\pi(x_1, \mathbf{F}_1) < 0\} + \phi^W \mathbf{1}\{a_1^\pi(x_1, \mathbf{F}_1) > 0\}) s_1 a_1^\pi(x_1, \mathbf{F}_1) | x_1, \mathbf{F}'_0].$$

Thus, the claimed property holds for $m = 1$ (recall that $x_1 \equiv x_1^\pi$).

The cases corresponding to $m = 2, \dots, N - 1$ can be dealt with by recursively applying a logic similar to the case $m = 1$. \square

Proof of Proposition 3 (Bounds on deltas). We derive the inequality $\Delta_{n,m}^\pi(t, x_n, \mathbf{G}_n(t)) \leq \bar{\delta}(t, T_m) \phi^W C^W$. The inequality $\Delta_{n,m}^\pi(t, x_n, \mathbf{G}_n(t)) \geq \bar{\delta}(t, T_m) \phi^I C^I$ can be established in an analogous manner. It holds that

$$\begin{aligned} \Delta_{n,m}^\pi(t, x_n, \mathbf{G}_n(t)) &= \frac{\bar{\delta}(t, T_m)}{F(t, T_m)} \mathbb{E} [(\phi^I A_m^\pi(x_m^\pi, \mathbf{F}_m) \mathbf{1}\{A_m^\pi(x_m^\pi, \mathbf{F}_m) < 0\} \\ &\quad + \phi^W A_m^\pi(x_m^\pi, \mathbf{F}_m) \mathbf{1}\{A_m^\pi(x_m^\pi, \mathbf{F}_m) > 0\}) s_m | x_n, \mathbf{G}_n(t)] \\ &\leq \frac{\bar{\delta}(t, T_m)}{F(t, T_m)} \mathbb{E} [\phi^W C^W s_m | x_n, \mathbf{G}_n(t)] \\ &= \bar{\delta}(t, T_m) \phi^W C^W \frac{\mathbb{E}[s_m | F(t, T_m)]}{F(t, T_m)} \\ &= \bar{\delta}(t, T_m) \phi^W C^W, \end{aligned}$$

where the first and last equalities, respectively, follow from Proposition 2 and the property $\mathbb{E}[s_m|F(t, T_m)] = F(t, T_m)$. \square

Proof of Lemma 1 (Characterization). The claimed properties are true in stage $N - 1$, as $V_{N-1}(x_{N-1}, s_{N-1}) = (\phi^W s_{N-1} - c^W)^+ (x_{N-1} \wedge C^W)$, and $\underline{b}_{N-1}(s_{N-1}) = 0$ and $\bar{b}_{N-1}(s_{N-1}) = \bar{x}1\{\phi^W s_{N-1} - c^W \leq 0\}$. Make the induction hypothesis that these properties also hold in stages $n + 1, \dots, N - 2$. Consider stage n . By definition, $W_n(x, \mathbf{F}'_n)$ is the discounted expected value of piecewise linear continuous functions, each with possible break points only in set \mathcal{Q} , which implies that this function satisfies the same property. The quantities $\underline{b}_n(\mathbf{F}_n)$ and $\bar{b}_n(\mathbf{F}_n)$ are optimal solutions to the following maximizations, respectively (see the proof of Secomandi 2010, Theorem 1):

$$\begin{aligned} & \max_{x_{n+1} \in \mathcal{X}} W_n(x_{n+1}, \mathbf{F}'_n) - (\phi^I s_n + c^I) x_{n+1}, \\ & \max_{x_{n+1} \in \mathcal{X}} W_n(x_{n+1}, \mathbf{F}'_n) - (\phi^W s_n - c^W) x_{n+1}. \end{aligned}$$

Thus, $\underline{b}_n(\mathbf{F}_n)$ and $\bar{b}_n(\mathbf{F}_n)$ can be taken to belong to set \mathcal{Q} . It follows that (recalling that $C^I < 0$)

$$V_n(x_n, \mathbf{F}_n) = \begin{cases} (\phi^I s_n + c^I) C^I + W_n(x_n - C^I, \mathbf{F}'_n), & x_n \in [0, \underline{b}_n(\mathbf{F}_n) + C^I], \\ (\phi^I s_{n-1} + c^I) [x_n - \underline{b}_n(\mathbf{F}_n)] + W_n(\underline{b}_n(\mathbf{F}_n), \mathbf{F}'_n), & x_n \in [\underline{b}_n(\mathbf{F}_n) + C^I, \underline{b}_n(\mathbf{F}_n)], \\ W_n(x_n, \mathbf{F}'_n), & x_n \in [\underline{b}_n(\mathbf{F}_n), \bar{b}_n(\mathbf{F}_n)], \\ (\phi^W s_{N-1} - c^W) [x_n - \bar{b}_n(\mathbf{F}_n)] + W_n(\bar{b}_n(\mathbf{F}_n), \mathbf{F}'_n), & x_n \in (\underline{b}_n(\mathbf{F}_n), \bar{b}_n(\mathbf{F}_n) + C^W], \\ (\phi^W s_{N-1} - c^W) C^W + W_n(x_n - C^W, \mathbf{F}'_n), & x_n \in (\bar{b}_n(\mathbf{F}_n) + C^W, \bar{x}]. \end{cases}$$

It is easy to check that this function is piecewise linear and continuous in x_n with break points in set \mathcal{Q} for each given \mathbf{F}_n . Therefore, the claimed properties are true in stage n . By the principle of mathematical induction, they hold in every stage. \square

Proof of Proposition 5 (Representation and Lipschitz continuity). (a) This part follows directly from part (b) of Lemma 1.

(b) The claimed property holds trivially in stage N with $L_N(x) = 0$ for all $x \in \mathcal{X}$. Make the induction hypothesis that this property holds in stages $n + 1, \dots, N - 1$. Consider stage n and pick $\mathbf{F}_n^1, \mathbf{F}_n^2 \in \mathfrak{R}_+^{N-n}$.

Define $C := |C^I| \vee C^W$ and $C' := (\phi^I \vee \phi^W)C$. For each $a \in [-C, 0)$, it holds that

$$\begin{aligned} |p(a, s_n^2) - p(a, s_n^1)| &= |(\phi^I s_n^2 + c^I)a - (\phi^I s_n^1 + c^I)a| \\ &\leq \phi^I |a| |s_n^2 - s_n^1| \leq \phi^I C |s_n^2 - s_n^1| \leq C' |s_n^2 - s_n^1|. \end{aligned}$$

For $a = 0$, it holds that $|p(a, s_n^2) - p(a, s_n^1)| = 0 \leq \phi^W C |s_n^2 - s_n^1| \leq C' |s_n^2 - s_n^1|$.

For each $a \in (0, C]$, it holds that

$$\begin{aligned} |p(a, s_n^2) - p(a, s_n^1)| &= |(\phi^W s_n^2 - c^W)a - (\phi^W s_n^1 - c^W)a| \\ &\leq \phi^W |a| |s_n^2 - s_n^1| \leq \phi^W C |s_n^2 - s_n^1| \leq C' |s_n^2 - s_n^1|. \end{aligned}$$

Thus, given $x \in \mathcal{X}$ and $a \in \mathcal{A}(x) \subseteq [-C, C]$, we have

$$|p(a, s_n^2) - p(a, s_n^1)| \leq C' |s_n^2 - s_n^1|. \quad (33)$$

Replacing $W_n(x_{n+1}, \mathbf{F}_n^{2'}) - W_n(x_{n+1}, \mathbf{F}_n^{1'})$ with $\check{W}_n(x_{n+1}, \mathbf{F}_n^{2'}, \mathbf{F}_n^{1'})$ for expositional convenience, it holds that

$$\begin{aligned} \left| \check{W}_n(x_{n+1}, \mathbf{F}_n^{2'}, \mathbf{F}_n^{1'}) \right| &= \delta \left| \mathbb{E} \left[V_{n+1}(x_{n+1}, \mathbf{F}_n^{2'} \vec{\beta}^{n, n+1}(\mathbf{Y})) - V_{n+1}(x_{n+1}, \mathbf{F}_n^{1'} \vec{\beta}^{n, n+1}(\mathbf{Y})) \right] \right| \\ &\leq \delta \mathbb{E} \left[\left| V_{n+1}(x_{n+1}, \mathbf{F}_n^{2'} \vec{\beta}^{n, n+1}(\mathbf{Y})) - V_{n+1}(x_{n+1}, \mathbf{F}_n^{1'} \vec{\beta}^{n, n+1}(\mathbf{Y})) \right| \right] \\ &\leq \delta \mathbb{E} \left[L_{n+1}(x_{n+1}) \sum_{m=n+1}^N |F_{n,m}^2 - F_{n,m}^1| \beta_m^{n, n+1}(\mathbf{Y}) \right] \\ &= \delta L_{n+1}(x_{n+1}) \sum_{m=n+1}^N |F_{n,m}^2 - F_{n,m}^1| \mathbb{E} [\beta_m^{n, n+1}(\mathbf{Y})] \\ &= \delta L_{n+1}(x_{n+1}) \sum_{m=n+1}^N |F_{n,m}^2 - F_{n,m}^1|, \end{aligned} \quad (34)$$

where the first inequality holds by Jensen's inequality, the second inequality follows from the induction hypothesis, and the last equality follows from (23).

Given $x \in \mathcal{X}$ and $a \in \mathcal{A}(x)$, inequalities (33) and (34) imply that for all $\mathbf{F}_n^1, \mathbf{F}_n^2 \in \mathfrak{R}_+^{N-n}$ it holds that

$$\begin{aligned} |v_n(x, a, \mathbf{F}_n^2) - v_n(x, a, \mathbf{F}_n^1)| &= \left| p(a, s_n^2) - p(a, s_n^1) + W_n(x_n - a, \mathbf{F}_n^{2'}) - W_n(x_n - a, \mathbf{F}_n^{1'}) \right| \\ &\leq |p(a, s_n^2) - p(a, s_n^1)| + \left| W_n(x_n - a, \mathbf{F}_n^{2'}) - W_n(x_n - a, \mathbf{F}_n^{1'}) \right| \\ &\leq C' |s_n^2 - s_n^1| + \delta L_{n+1}(x_n - a) \sum_{m=n+1}^N |F_{n,m}^2 - F_{n,m}^1| \\ &\leq \{C' \vee [\delta L_{n+1}(x_n - a)]\} \sum_{m=n}^N |F_{n,m}^2 - F_{n,m}^1|. \end{aligned}$$

Thus, the function $v_n(x, a, \mathbf{F}_n)$ is Lipschitz continuous in $\mathbf{F}_n \in \mathfrak{R}_+^{N-n}$ for each given $x \in \mathcal{X}$ and $a \in \mathcal{A}(x)$.

Part (a) of this proposition and Dudley (2002, Proposition 11.2.2(a), p. 391) imply that the claimed property holds in stage n . It follows from the principle of mathematical induction that the claimed property holds in every stage. \square

Lemma 3. Lemma 3 is needed in the proof of Proposition 6.

Lemma 3 (Differentiability and unique optimal action). *Under Assumption 2, for every $n \in \mathcal{N}$, if $V_n(x_n, \mathbf{F}_n)$ is differentiable with respect to each element of \mathbf{F}_n at a given futures curve $\bar{\mathbf{F}}_n$ for given $x_n \in \mathcal{X}$, then at $(x_n, \bar{\mathbf{F}}_n)$ there is a unique optimal action, denoted as $a_n^*(x_n, \bar{\mathbf{F}}_n)$, and*

$$\left. \frac{\partial V_n(x_n, \mathbf{F}_n)}{\partial F_{n,m}} \right|_{\mathbf{F}_n = \bar{\mathbf{F}}_n} = \left. \frac{\partial v_n(x_n, a_n^*(x_n, \bar{\mathbf{F}}_n), \mathbf{F}_n)}{\partial F_{n,m}} \right|_{\mathbf{F}_n = \bar{\mathbf{F}}_n}, \quad \forall m \in \mathcal{N}_n.$$

Proof. Consider stage n . Part (b) of Proposition 5 and Rademacher's theorem imply that the function $V_n(x_n, \mathbf{F}_n)$ is differentiable almost everywhere in each element of \mathbf{F}_n for each given $x_n \in \mathcal{X}$. Fix x_n and pick $\bar{\mathbf{F}}_n$ such that $V_n(x_n, \mathbf{F}_n)$ is differentiable in each element of \mathbf{F}_n at $\bar{\mathbf{F}}_n$. In particular, this means that $\partial V_n(x_n, \mathbf{F}_n)/\partial s_n$ exists at \bar{s}_n . Each function $v_n(x_n, a, \mathbf{F}_n)$ with $a \in \mathcal{A}'(x_n)$ is linear in s_n . This and part (a) of Proposition 5 imply that there is a unique optimal action at $\bar{\mathbf{F}}_n$, because otherwise $\partial V_n(x_n, \mathbf{F}_n)/\partial s_n$ would not exist at \bar{s}_n . This implies that $V_n(x_n, \bar{\mathbf{F}}_n) \equiv v_n(x_n, a_n^*(x_n, \bar{\mathbf{F}}_n), \bar{\mathbf{F}}_n)$ in a neighborhood of \mathbf{F}_n . The assumed differentiability of $V_n(x_n, \mathbf{F}_n)$ in each element of \mathbf{F}_n at $\bar{\mathbf{F}}_n$ implies that $\partial V_n(x_n, \mathbf{F}_n)/\partial F_{n,m} = \partial v_n(x_n, a_n^*(x_n, \bar{\mathbf{F}}_n), \mathbf{F}_n)/\partial F_{n,m}$ at $\bar{F}_{n,m}$ for all $m \in \mathcal{N}_n$. \square

Proof of Proposition 6 (Pathwise deltas with an optimal policy). Under Assumption 2, an optimal policy satisfies the conditions (a) and (b) stated in Assumption 1 by Proposition 5(b) and Lemma 3, respectively. Expression (22) then follows from Proposition 2. \square

References

- Broadie, M., P. Glasserman. 1996. Estimating security price derivatives using simulation. *Management Science* **42**(2) 269–285.
- Dudley, R. M. 2002. *Real Analysis and Probability*. 2nd ed. Cambridge University Press, Cambridge, UK.
- Hull, J. C. 2012. *Options, Futures, and Other Derivatives Securities*. Eight ed. Prentice Hall, Englewood Cliffs, NJ, USA.
- 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.
- Luenberger, D. G. 1998. *Investment Science*. Oxford University Press, New York, NY, USA.
- Secomandi, N. 2010. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.
- Smith, J. E., K. F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research* **47**(1) 1–15.