

Commodity Procurement with Demand Forecast and Forward Price Updates

Nicola Secomandi, Sunder Kekre

Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA
15213-3890, USA, {ns7, sk0a}@andrew.cmu.edu

Tepper Working Paper 2009-E25

September 2009

Abstract

Commodities, ranging from natural gas to memory chips, can be procured both by trading on the date in spot markets and in advance in forward markets. Transaction costs, such as brokerage fees, are typically higher in spot markets than in forward markets. Moreover, the forecast of a firm's commodity requirement (demand) for a given future date typically changes in an uncertain fashion over time. Thus, although the dynamics of forward and spot prices are notoriously uncertain, firms that procure commodities face the dilemma of choosing between early and possibly less expensive commitments with residual demand uncertainty and late and possibly more expensive sourcing of the exact amount needed. We investigate this issue by developing and analyzing a model of commodity procurement for a single future date. Our model generalizes models available in the real options and operations management literature, by simultaneously considering correlated demand forecast and forward price updates in a setting characterized by multiple forward transactions and a single spot transaction. We derive the structure of the optimal procurement policy and discuss its computation in cases of practical interest. In a numerical study, based on applying our model to natural gas data, we offer managerial insights on the effects that demand forecast and forward price updates, both in isolation and combined, have on the value of a firm's procurement policy. We also assess the sensitivities of these effects to parameters of interest and the potential managerial relevance of the combined effect. Our model and results have significance beyond the specific application.

1. Introduction

Commodities play a key economic role in supply chains, as they constitute primary inputs to most industrial and commercial activities (Geman [20]). They range from energy sources, such as electricity and natural gas, to raw materials and components, such as metals, agricultural products, and memory chips. They are traded both in spot markets for immediate delivery and in forward markets for future delivery. Their procurement for further processing, distribution, or usage is an important problem in practice.

The basic tradeoff in commodity procurement is as follows. Spot and forward prices are notoriously uncertain and volatile (Geman [20]), but it is well known that transaction costs, such as brokerage fees, incurred by commodity procurement firms to access these markets are higher in spot than in forward markets (Williams [65, 66]). The commodity requirement (demand) is itself not known exactly by a firm until the usage date, with the forecast of this amount evolving in an uncertain fashion over time. Thus, the dilemma faced by commodity procurement firms is how much to procure in advance in the forward market, possibly at a cheaper price but also with

residual demand uncertainty, and on the date in the spot market, possibly at a higher price but with exact knowledge of the demand.

Companies have taken different approaches to deal with this dilemma. One is that of an energy reselling company, with which we recently collaborated and that motivated this research. This company purchases energy, such as natural gas and electricity, and resells it to satisfy the demand of commercial and industrial customers. The operating mode of this company can be roughly summarized as the following two step approach: (1) At some hypothetical initial time, such as the time when a new customer is added, procure by trading in the forward market an amount of commodity equal to the (deterministic) demand forecast in each of the months included in the planning horizon; (2) as the first month in the planning horizon becomes the current month, trade in the forward market to secure an amount of supply equal to the demand forecast for the month added to the planning horizon, and trade in the spot market to make up for any differences between the realized demand and the amount procured earlier in the forward market for the current month. (This is clearly a convenient expositional simplification since demand for the month does not become known at once at the beginning of the month.) This approach does not involve optimization of the procurement choices, nor updating of the supply procured in the forward market until the spot date.

Other energy reselling firms, local distribution companies (utilities), and industrial and commercial users have developed optimization based approaches to support their procurement choices. Examples include the stochastic programming approach used by The Peoples Gas Light and Coke Company to plan its gas supply (Knowles and Wirick [40]; see also the models and applications discussed by Guldman [23], Morris et al. [47], Bopp et al. [7], Butler and Dyer [9], and Guldman and Wang [24]), and the deterministic optimization approach described by Avery et al. [1] and applied at Southwest Gas Corporation and Questar Pipeline Corporation. Moreover, these companies appear to periodically reevaluate their supply plans, and hence their positions in the forward market, typically through reoptimization of the relevant models with updated parameters. For example, Levary and Dean [45] propose an optimization model that periodically updates its demand forecast in a reactive fashion and revises its earlier supply decisions accordingly; these authors discuss an application of their model to data pertaining to the East Ohio Gas Company.

This optimization based approach with periodic reevaluation of the supply plan contrasts the simple approach taken by the stated energy reselling firm. The natural question that arises is to what extent information updating can impact the financial performance of a commodity procurement policy. This issue is not well understood in the literature (reviewed in §2). The cited

optimization based approaches used in practice are rich in their contextual details. It is challenging to recreate all the details of the models and the data used in the specific applications. Thus, we investigate the stated issue by developing and analyzing a dynamic model of commodity procurement, which captures the salient features of this business process: differential transaction costs between forward and spot trading and demand forecast and forward price updates.

Specifically, we develop a Markov decision process (MDP; Stokey and Lucas [61], Puterman [50], Bertsekas [4]) where a firm procures a commodity for usage at a single future date. This firm can trade multiple times in the forward market, in advance of the usage date. Trading in the forward market may be advantageous because it commands lower transaction costs than trading in the spot market (we rule out arbitrage opportunities from trading in the forward and spot markets). The commodity procured in the forward market is delivered to the firm on the usage date, and the firm satisfies demand by making up any supply underage or overage by trading in the spot market. At each trading time, the firm has available newly updated information about the demand forecast and the forward price. We model the evolution of these two types of information as correlated Markov processes.

We establish that the optimal trading policy is of the basestock type, with two buy-up-to and sell-down-to levels that depend on the stage and the prevailing forward price and demand forecast at that time. In other words, in a given stage and for a given information set, it is optimal to increase (respectively, decrease) the total supply so far procured in the forward market to the buy-up-to (respectively, sell-down-to) basestock level if it is below (respectively, above) this level, and do nothing otherwise. Moreover, we show that in cases of practical interest these basestock levels can only take on a finite number of values that can be easily determined a priori. Hence, their computation reduces from solving an infinite dimensional problem to solving a finite dimensional problem. These characterizations are significant because they greatly simplify the computation of an optimal policy.

We apply our model to natural gas data and other data available in the literature in a computational study, where the joint evolution of the demand forecast and forward price update processes is based on variants of the models of Hausman [27] and Black [6]. We compare the value of an optimal policy to our model to those of simpler basestock type policies, whose basestock levels in each stage depend either on forward price or demand forecast updates, but not both, and of static policies that do not consider information updates (these are basestock policies that can trade in the forward market only at the beginning of the time horizon). These comparisons allow us to measure the effect of incorporating demand forecast, forward price, and both demand forecast and forward

price updates in a commodity procurement model that does not use such information, and may or may not optimize its initial supply decision.

We find that it is nontrivial to incorporate demand forecast updates in a firm's commodity procurement policy, because the demand forecast effect is negative. This means that price information and optimization are key elements for exploiting demand forecast updates in commodity procurement. Even though ignoring price information is detrimental in making commodity procurement choices, the forward price update effect is favorable but small. The combined forward price and demand forecast update effect is also positive but substantially larger than the forward price update effect. These findings suggest a complementary relationship between these two types of updates, and a substantially larger benefit from dynamic optimization when using both demand forecast and forward price updates than only forward price updates.

Increased spot demand variability (coefficient of variation, CV), due to either increased demand forecast volatility or length of the time horizon, reinforces these findings, which are instead largely insensitive to the forward price volatility. The instantaneous correlation between the demand forecast and the forward price updates impacts the demand forecast effect and the forward price effect, but does not substantially impact the combined effect.

These results are generic in the sense that they do not depend on the specific values assigned to the initial demand forecast and forward price of natural gas employed in our numerical study. In other words, the percent improvements in the values of the commodity procurement policies taken as benchmarks that we discuss in §5 do not depend on these initial values (this is due to how we model the demand forecast and forward price evolutions).

Although these percent improvements are small, being no more than 2%, their corresponding absolute improvements (cost savings) range from \$0.3M to \$1.7M per month (M denotes million). This suggests that fine tuning a firm's commodity procurement policy by jointly leveraging demand forecast and forward price updates may have potential managerial relevance. For example, these cost savings directly translate to margin improvements for unregulated natural gas resellers; given that these companies operate with very low margins, these margin improvements are likely to be substantial. For those local distribution companies that operate under benchmark based regulation, as discussed in §5, a fraction of these cost savings would translate to margin improvements, with their balance benefiting consumers. Different from the discussed percent improvements, these absolute improvements are specific to the application domain. Nonetheless, we believe that they should remain potentially relevant in other industrial and commercial settings, such as food processing, metal and chemical manufacturing, and the management of large restaurant and hotel chains.

The remainder of this paper is organized as follows. We review the related literature in §2. We present our model in §3. We derive the structure of its optimal policy in §4. We conduct our numerical study in §5. We conclude in §6.

2. Literature Review

Our work is related to the real options and the operations management literatures that deal with long- and short-term contracting in business-to-business settings. Kleindorfer and Wu [39] and Kleindorfer [38] review this literature.

The real options literature (Sick [57], Trigeorgis [64]) dealing with commodity industries emphasizes the modeling of uncertain price evolution (Seppi [56]). The basic models of procurement in this literature are the swing option (see, e.g., Jaillet et al. [33]), a contract that gives its owner the ability to purchase a given amount of commodity a fixed number of times at a fixed price during a given time period, and other contracts with varying amounts of sourcing flexibility (see, e.g., Kamrad and Ritchken [37], Li and Kouvelis [46], Dong and Hong [14]). The operations management literature emphasizes the modeling of the uncertain evolution of the demand forecast in procurement. Basic models of demand forecast evolution in this literature are those proposed by Hausman [27] and Hausman and Peterson [28], and later extended and refined by Graves et al. [22] and Heath and Jackson [29], and models based on Bayesian updates, such as that of Eppen and Iyer [16]. Thus, the cited real options works model the uncertainty in the purchase price evolution, but neglect demand forecast uncertainty, whereas the cited operations management works do the opposite. In contrast, our model captures the uncertain evolution of both the forward price and the demand forecast.

There is a growing body of work, mainly in the operations management and operations research literatures, on commodity procurement with both purchase price and demand uncertainty, including Kalyon [35], an early reference, Gavirneni [19], Seifert et al. [55], Berling and Rosling [3], Berling and Martínez-de-Albéniz [2], Boyabatli et al. [8], Goel and Gutierrez [21], Nascimento and Powell [49], and Yang and Xia [67]. These models assume that demand is independent of the stochastic purchase price and do not account for the evolution of the demand forecast, whereas we propose a procurement model with correlated forward (purchase) price and demand forecast updates.

Gurnani and Tang [25] study a two period model with independent purchase price and demand forecast updates; specifically, the demand update depends on a market signal, which changes over time, but the price does not depend on this signal. Gaur and Seshadri [17] consider an inventory model where demand is correlated with the price of a financial asset, but the purchase price is

deterministic. Unlike these models, the stochastic processes that represent the forward price and demand forecast updates are correlated in our model.

Gaur et al. [18] consider a finite horizon inventory model with correlated demand forecast and selling price updates, deterministic and time dependent purchase price, and a single purchase opportunity to be chosen within the time horizon. Different from their model, our model features correlated demand forecast and purchase, rather than selling, price, and multiple sourcing opportunities. Ritchken and Tapiero [51] consider a two period inventory model with correlated purchase price and demand uncertainty, as well as options on inventory purchases in the second period, which maximizes the expected utility of a risk averse manager. In contrast, our model features more than two periods, supply is obtained through purchases in the forward and spot markets, no inventory is held over time, and its objective is not expected utility maximization.

Our analysis of the optimal policy is related to those of Secomandi [54] and Lai et al. [43]. However, our model includes a demand component that is absent from their models.

3. Model

There is a finite horizon. The firm satisfies a requirement for a commodity at the end of this time horizon. We refer to this requirement as the spot demand and denote it by d . This quantity is a random variable as of any time before the end of the time horizon. The firm can satisfy this requirement by procuring the commodity in the spot market at the end of the time horizon. The firm can also procure supply in advance by entering into forward contracts with physical delivery at the end of the time horizon. We subdivide this time horizon into $J + 1$ discrete time periods in set $\mathcal{J} := \{1, \dots, J + 1\}$.

We denote by F_j the forward price at the beginning of period $j \in \mathcal{J}$. The forward price evolves during the time horizon as a known stochastic process, which is assumed to satisfy the Markov property. At the end of the time horizon, the forward price is equal to the spot price; that is, $F_{J+1} \equiv s$, where s denotes the spot price of the commodity at this time.

The firm trades in spot and forward markets through a broker and pays brokerage fees (transaction costs) to this broker to access these markets. A well known feature of brokerage fees is that those associated with the spot market are typically larger than those associated with the forward market. This is due to the broker having to arrange spot transactions in a hurry. Williams [65, p. 146] and [66] discusses this aspect in detail.

We model this feature as follows. If the firm purchases spot one unit of commodity from the broker, it pays the broker an amount $(1 + A)s$; if the firm sells spot one unit of commodity to the

broker, it receives from the broker an amount $(1 - A)s$, with $A \in (0, 1)$. At the beginning of each period j , it can forward purchase one unit of commodity from the broker at price $(1 + B)F_j$, and it can forward sell one unit of commodity to the broker at price $(1 - B)F_j$, with $B \in (0, 1)$; in the latter case, we only allow the firm to forward sell commodity that it previously purchased, that is, we do not allow the firm to short sell the commodity. Consistent with what stated above, we assume that $B < A$.

Our modeling of transaction costs as proportional is consistent with models studied in the finance literature (see, e.g., Constantinides et al. [10]). However our analysis also applies to the case of additive transaction costs, whereby the price received and paid from selling and purchasing one unit of the commodity spot are $s - A$ and $s + A$, respectively (analogous expressions would hold when trading in the forward market).

As the forward price, the firm's demand forecast also evolves over time. Specifically, at the beginning of each period j , the firm has available a forecast D_j of its spot demand d . Analogous to the forward price process, this demand forecast evolves during the time horizon as a known stochastic process, which satisfies the Markov property and is equal to the spot demand at the end of this time horizon ($D_{J+1} \equiv d$). The price and demand forecast stochastic processes are assumed to be correlated.

For notational convenience, we define $\mathbb{E}_j[\cdot] := \mathbb{E}[\cdot | D_j, F_j]$, $\forall j \in \mathcal{J} \setminus \{J + 1\}$, $(D_j, F_j) \in \mathfrak{R}_+^2$, expectation with respect to a suitable distribution, discussed at the beginning of §4, given that D_j and F_j are known in this stage. Notice that $\mathbb{E}_j[\cdot]$ is not a random variable in stage j because D_j and F_j are known in this stage. That is, our notation does not distinguish between random variables and their realizations; which is which should be clear from the context. Alternatively, one would have to introduce additional notation for random variables, such as \tilde{D}_j and \tilde{F}_j to denote the random variables demand forecast and forward price in stage j . In this case, it would hold that $\mathbb{E}_j[\cdot] \equiv \mathbb{E}[\cdot | \tilde{D}_j = D_j, \tilde{F}_j = F_j]$. We do not employ this notational style in this paper for expositional simplicity.

The problem faced by the firm is to decide how to procure supply by trading in the forward market during the finite time horizon to meet the demand faced at the end of this horizon. Notice that the firm will always provide 100% service level to its customers, as it must satisfy the entire demand that it faces. Indeed, a feasible procurement policy is to simply wait till the end of the time horizon, observe the realized demand, and procure the corresponding amount of commodity by trading on the spot market. However, this may not be optimal.

Thus, we formulate the procurement problem as an MDP. The stages are those in set \mathcal{J} . In

each stage j , the state is $(x_j, D_j, F_j) \in \mathfrak{R}_+^3$, where x_j is the amount of supply procured in the forward market up to stage j . The optimal value function in stage j and state (x_j, D_j, F_j) is denoted as $V_j(x_j, D_j, F_j)$, and is expressed in stage $J + 1$ dollars (because money between the firm and the borker is exchanged only at this time). It is useful to define $(\cdot)^+ := \max\{\cdot, 0\}$ and $(\cdot)^- := -\min\{\cdot, 0\}$. The MDP formulation follows

$$V_{J+1}(x, d, s) := (1 - A)s(x - d)^+ - (1 + A)s(x - d)^-, \quad \forall (x, d, s) \in \mathfrak{R}_+^3 \quad (1)$$

$$\begin{aligned} V_j(x_j, D_j, F_j) &= \max_{a \leq x_j} (1 - B)F_j a^+ - (1 + B)F_j a^- + \mathbb{E}_j[V_{j+1}(x_j - a, D_{j+1}, F_{j+1})], \\ &\quad \forall (x_j, D_j, F_j) \in \mathfrak{R}_+^3. \end{aligned} \quad (2)$$

The first term in (1) is the revenue collected from selling excess contracted supply on the spot market, and the second term is the cost of any supply shortfall. The first term in (2) is the revenue collected in stage j and state (x_j, D_j, F_j) by forward selling an amount a^+ of *previously* bought supply ($a \leq x_j$), the second is the cost of forward purchasing additional supply a^- , and the third is the expected value of acting optimally from the resulting state in stage $j + 1$.

4. Optimal Policy Analysis

This section analyzes the structure of an optimal policy by imposing more structure on model (1)-(2).

We first discuss the distribution associated with the definition of \mathbb{E}_j . Model (1)-(2) can be interpreted in different ways. In one interpretation, the firm maximizes expected margin, and the distribution in question reflects the firm's own assessment of the probabilistic evolution of the forward price and its demand forecast. In another interpretation, the firm maximizes the market value of its operations and, under some assumptions, model (1)-(2) can be justified by the application of risk neutral valuation theory; that is, the distribution in question can be taken to be the so called risk neutral distribution of these random variables, possibly obtained using equilibrium arguments (see, e.g., Cox et al. [12], Duffie [15], Hull [32, Chapters 25, 31]; papers that have applied risk neutral valuation to operational problems include Singhal [58], Triantis and Hodder [63], Lederer and Singhal [44], Kogut and Kulatilaka [41], Huchzermeier and Cohen [31], Kamrad and Lele [36], Kouvelis [42], Smith and McCardle [60], Birge [5], and Gaur et al. [18]). When not all of the assumptions required to apply this approach are satisfied, our model formulation is consistent with the heuristic application of this theory, as discussed, for example, by Smith [59, p. 95] (see also Schwartz [53]). We discuss these issues in more detail in Appendix A, since they

are not central to our analysis. We proceed by making a few assumptions, which are consistent with either one of these interpretations.

Assumption 1 states that the forward price process is a martingale. In the expected margin interpretation, this means that the firm is not in the business of trading forward contracts with a speculative motive. In the risk neutral interpretation, this is a natural assumption when the risk free rate is assumed to be deterministic and constant, since this implies that the forward price of the commodity is equal to its futures price (Cox et al. [11]), which evolves as a martingale in the risk neutral world (see, e.g., Duffie [15]).

Assumption 1 (Forward price is a martingale). It holds that $\mathbb{E}[F_{j'}|F_j] = F_j$, $\forall j \in \mathcal{J} \setminus \{J+1\}$, $j' \in \mathcal{J}$, $j' > j$, and $F_j \in \mathfrak{R}_+$.

This assumption implies that the expectations of a future forward price and the spot price conditional on the current forward price are finite. Assumption 2 further states that other relevant conditional expectations are finite.

Assumption 2 (Bounded conditional means). It holds that $\mathbb{E}[d|D_j] < \infty$, $\forall j \in \mathcal{J}$ and $D_j \in \mathfrak{R}_+$, and $\mathbb{E}[sd|D_j, F_j] < \infty$, $\forall j \in \mathcal{J}$ and $(D_j, F_j) \in \mathfrak{R}_+^2$.

For the ensuing analysis, it is convenient to define the optimal continuation function as

$$\begin{aligned} U_{J+1}(x, d, s) &:= 0, \forall (x, d, s) \in \mathfrak{R}_+^3 \\ U_j(x_j, D_j, F_j) &:= \mathbb{E}[V_{j+1}(x_j, D_{j+1}, F_{j+1})|D_j, F_j], \forall (x_j, D_j, F_j) \in \mathfrak{R}_+^3. \end{aligned}$$

Proposition 1 bounds the optimal value and continuation functions, and shows that they are both finite. These properties are used in Proposition 2.

Proposition 1 (Bounds and finiteness). *Suppose that Assumptions 1-2 hold. The functions $V_j(x_j, D_j, F_j)$ and $U_j(x_j, D_j, F_j)$ are bounded below by $-(1+A)\mathbb{E}_j[sd] - (1-A)F_jx_j$ and above by $-(1-B)\mathbb{E}_j[sd] + (1-B)F_jx_j$, $\forall (x_j, D_j, F_j) \in \mathfrak{R}_+^3$. Moreover, it holds that $|U_j(x_j, D_j, F_j)| < \infty$ and $|V_j(x_j, D_j, F_j)| < \infty$, $\forall j \in \mathcal{J}$ and $(x_j, D_j, F_j) \in \mathfrak{R}_+^3$.*

Proof. Consider stage $J+1$. It is evident that $|V_{J+1}(x, d, s)| < \infty$ and $|U_{J+1}(x, d, s)| < \infty$, $\forall (x, d, s) \in \mathfrak{R}_+^3$. For every state $(x, d, s) \in \mathfrak{R}_+^3$, it holds that

$$\begin{aligned} V_{J+1}(x, d, s) &= (1-A)s(x-d)^+ - (1+A)s(x-d)^- \\ &\geq (1-A)s(-x) + (1+A)s(-d) = -(1+A)sd - (1-A)sx, \end{aligned}$$

and

$$\begin{aligned} V_{J+1}(x, d, s) &= (1 - A)s(x - d)^+ - (1 + A)s(x - d)^- \\ &\leq (1 - B)s(x - d)^+ - (1 - B)s(x - d)^- = -(1 - B)sd + (1 - B)sx. \end{aligned}$$

Since $sd \equiv \mathbb{E}_{J+1}[sd]$ and $s \equiv \mathbb{E}_{J+1}[s] \equiv F_{J+1}$, the claimed properties hold in stage $J + 1$. Make the induction hypothesis that they also hold in stages $j + 1, \dots, J$. In particular, this means that for all $(x, D_{j+1}, F_{j+1}) \in \mathfrak{R}_+^3$ it holds that

$$-(1 + A)\mathbb{E}_{j+1}[sd] - (1 - A)F_{j+1}x \leq V_{j+1}(x, D_{j+1}, F_{j+1}) \leq -(1 - B)\mathbb{E}_{j+1}[sd] + (1 - B)F_{j+1}x.$$

Consider stage j . Pick arbitrary and feasible state (x_j, D_j, F_j) . The induction hypothesis and Assumption 1 imply that

$$\begin{aligned} U_j(x_j, D_j, F_j) &\equiv \mathbb{E}_j[V_{j+1}(x_j, D_{j+1}, F_{j+1})] \\ &\geq \mathbb{E}_j[-(1 + A)\mathbb{E}_{j+1}[sd] - (1 - A)F_{j+1}x] \\ &= -(1 + A)\mathbb{E}_j[sd] - (1 - A)F_jx, \end{aligned}$$

and

$$\begin{aligned} U_j(x_j, D_j, F_j) &\equiv \mathbb{E}_j[V_{j+1}(x_j, D_{j+1}, F_{j+1})] \\ &\leq \mathbb{E}_j[-(1 - B)\mathbb{E}_{j+1}[sd] + (1 - B)F_{j+1}x] \\ &= -(1 - B)\mathbb{E}_j[sd] + (1 - B)F_jx. \end{aligned}$$

Notice that

$$\begin{aligned} V_j(x_j, D_j, F_j) &\leq \max_{a \leq x_j} (1 - B)F_ja + U_j(x_j - a, D_j, F_j) \\ &= (1 - B)F_jx_j + \max_{x_{j+1} \geq 0} U_j(x_{j+1}, D_j, F_j) - (1 - B)F_jx_{j+1}, \end{aligned}$$

where the equality holds by letting $x_{j+1} = x_j - a$. An optimal solution to this maximization does not depend on x_j ; denote it by $x_{j+1}^*(D_j, F_j)$. It follows that

$$\begin{aligned} V_j(x_j, D_j, F_j) &\leq (1 - B)F_jx_j + U_j(x_{j+1}^*(D_j, F_j), D_j, F_j) - (1 - B)F_jx_{j+1}^*(D_j, F_j) \\ &\leq (1 - B)F_jx_j - (1 - B)\mathbb{E}_j[sd] + (1 - B)F_jx_{j+1}^*(D_j, F_j) - (1 - B)F_jx_{j+1}^*(D_j, F_j) \\ &= -(1 - B)\mathbb{E}_j[sd] + (1 - B)F_jx_j. \end{aligned}$$

Moreover, doing nothing in stages j through J and transacting in the spot market in stage $J + 1$ is a feasible policy. Hence, it holds that

$$\begin{aligned} V_j(x_j, D_j, F_j) &\geq \mathbb{E}_j[(1 - A)s(x_j - d)^+ - (1 + A)s(x_j - d)^-] \\ &= -(1 + A)\mathbb{E}_j[sd] - (1 - A)F_jx_j. \end{aligned}$$

Assumption 2 implies that $|U_j(x_j, D_j, F_j)| < \infty$ and $|V_j(x_j, D_j, F_j)| < \infty$. Thus, the claimed properties hold in stage j . The principle of mathematical induction implies that they hold in every stage. \square

It is useful to define the feasible inventory and action set as $\mathcal{C} := \{(x, a) : x \in \mathfrak{R}_+, a \in (-\infty, x]\}$. It is easy to verify that this set is convex and a lattice. Proposition 2 establishes a useful property of the optimal value and continuation functions, which is then used to establish Proposition 3. (See Rockafellar [52, p. 24] for the definition of proper concave function.)

Proposition 2 (Proper concavity). *In each stage j , the functions $V_j(x_j, D_j, F_j)$ and $U_j(x_j, D_j, F_j)$ are proper concave in $x_j \in \mathfrak{R}_+$ for each given $(F_j, D_j) \in \mathfrak{R}_+^2$.*

Proof. It is clear that the claimed property holds in stage $J + 1$. Make the induction hypothesis that this property also holds in stages $j + 1$ through J . Consider stage j and the feasible pair (D_j, F_j) . Choose $\phi \in [0, 1]$ and two feasible inventory levels x^1 and x^2 . Define inventory level $x^\phi := \phi x^1 + (1 - \phi)x^2$, which is obviously feasible. By the induction hypothesis, it holds that

$$V_{j+1}(x^\phi, D_{j+1}, F_{j+1}) \geq \phi V_{j+1}(x^1, D_{j+1}, F_{j+1}) + (1 - \phi)V_{j+1}(x^2, D_{j+1}, F_{j+1}).$$

It follows that

$$U_j(x^\phi, D_j, F_j) \geq \phi U_j(x^1, D_j, F_j) + (1 - \phi)U_j(x^2, D_j, F_j),$$

which, together with Proposition 1, establishes the proper concavity of $U_j(\cdot, D_j, F_j)$ in its first argument given (D_j, F_j) .

Consider the feasible actions a^1 and a^2 at inventory level x^1 and x^2 , respectively. Define $a^\phi := \phi a^1 + (1 - \phi)a^2$. Notice that $(x^\phi, a^\phi) \in \mathcal{C}$, because \mathcal{C} is convex. It follows from the proper concavity of $U_j(\cdot, D_j, F_j)$ in its first argument for each given (D_j, F_j) that

$$U_j(x^\phi - a^\phi, D_j, F_j) \geq \phi U_j(x^1 - a^1, D_j, F_j) + (1 - \phi)U_j(x^2 - a^2, D_j, F_j),$$

so that $U_j(x - a, D_j, F_j)$ is jointly proper concave in $(x, a) \in \mathcal{C}$ for each given (D_j, F_j) . This and the proper concavity of $(1 - B)F_j a^+ - (1 + B)F_j a^-$ in a imply that the function $(1 - B)F_j a^+ - (1 + B)F_j a^- + U_j(x - a, D_j, F_j)$ is jointly proper concave in $(x, a) \in \mathcal{C}$ for each given (D_j, F_j) . Proposition B-4 in Heyman and Sobel [30, p. 525] implies that $V_j(x, D_j, F_j)$ is proper concave in $x \in \mathfrak{R}_+$ for each given (D_j, F_j) . Hence, the claimed property holds in stage j . The principle of mathematical induction implies that this property holds in every stage. \square

Proposition 3 establishes that the optimal policy has a double basestock level structure. Its proof is related to the proof of Theorem 1 in Lai et al. [43].

Proposition 3 (Double basestock optimal policy). *The optimal policy is characterized by two basestock levels $\underline{b}_j(D_j, F_j) \leq \bar{b}_j(D_j, F_j)$ that depend on the stage j and the pair (F_j, D_j) , such that in stage j and state (x_j, F_j, D_j) an optimal action, denoted as $a_j^*(x_j, F_j, D_j)$, is as follows:*

$$a_j^*(x_j, F_j, D_j) = \begin{cases} x_j - \underline{b}_j(D_j, F_j) & x \in [0, \underline{b}_j(D_j, F_j)) \\ 0 & x \in [\underline{b}_j(D_j, F_j), \bar{b}_j(D_j, F_j)] \\ x_j - \bar{b}_j(D_j, F_j) & x \in (\bar{b}_j(D_j, F_j), \infty) \end{cases} .$$

Proof. Consider arbitrary stage j and feasible state (x, D_j, F_j) . Proposition 2 and Lemma 2.6.2(b) in Topkis [62] imply that, given (D_j, F_j) , the function $U_j(x-a, D_j, F_j)$ is supermodular in $(x, a) \in \mathcal{C}$. Given (D_j, F_j) , the function $(1-B)F_j a^+ - (1+B)F_j a^-$ is trivially supermodular in $(x, a) \in \mathcal{C}$. It follows from Lemma 2.6.1(b) in Topkis [62] that, given (D_j, F_j) , the function $(1-B)F_j a^+ - (1+B)F_j a^- + U_j(x-a, D_j, F_j)$ is supermodular in $(x, a) \in \mathcal{C}$. Theorem 2.8.2 in Topkis [62] implies that, given (D_j, F_j) , any optimal action in stage j and state (x, D_j, F_j) increases in x . This establishes that there exist two feasible inventory levels $\underline{b}_j(D_j, F_j) \leq \bar{b}_j(D_j, F_j)$ that subdivide the feasible inventory set into the sets $[0, \underline{b}_j(D_j, F_j))$, $[\underline{b}_j(D_j, F_j), \bar{b}_j(D_j, F_j)]$, and $(\bar{b}_j(D_j, F_j), \infty)$, where buying, doing nothing, and selling is optimal, respectively.

Suppose that $x \in [0, \underline{b}_j(D_j, F_j))$. Analogous to the proof of Proposition 1, by letting $x_{j+1} = x - a$ and ignoring the constraint $x_{j+1} \geq x$ the relevant maximization can be written as

$$(1+B)F_j x + \max_{x_{j+1} \geq 0} U_j(x_{j+1}, D_j, F_j) - (1+B)F_j x_{j+1} .$$

This shows that an optimal solution to this maximization does not depend on x . Thus, $\underline{b}_j(D_j, F_j)$ is an optimal solution to this maximization, and hence $x - \underline{b}_j(D_j, F_j)$ is an optimal action at x . One can show that $x - \bar{b}_j(D_j, F_j)$ is an optimal action for $x \in (\bar{b}_j(D_j, F_j), \infty)$. Since 0 is an optimal action for $x \in [\underline{b}_j(D_j, F_j), \bar{b}_j(D_j, F_j)]$, the proof is complete. \square .

Corollary 1 follows immediately from Proposition 3. It is useful for establishing Proposition 4.

Corollary 1 (Characterization). *In every stage $j \in \mathcal{J} \setminus \{J+1\}$ it holds that*

$$V_j(x, F_j, D_j) = \begin{cases} (1+B)F_j x_j + \underline{K}_j(D_j, F_j) & x_j \in [0, \underline{b}_j(D_j, F_j)) \\ U_j(x_j, D_j, F_j) & x_j \in [\underline{b}_j(D_j, F_j), \bar{b}_j(D_j, F_j)] \\ (1-B)F_j x_j + \bar{K}_j(D_j, F_j) & x_j \in (\bar{b}_j(D_j, F_j), \infty) \end{cases} ,$$

where

$$\begin{aligned} \underline{K}_j(D_j, F_j) &:= U_j(\underline{b}_j(D_j, F_j), D_j, F_j) - (1+B)F_j \underline{b}_j(D_j, F_j) \\ \bar{K}_j(D_j, F_j) &:= U_j(\bar{b}_j(D_j, F_j), D_j, F_j) - (1-B)F_j \bar{b}_j(D_j, F_j) . \end{aligned}$$

Computation of the optimal basestock levels in each stage and for each demand forecast and forward price pair is greatly facilitated by imposing additional structure on the demand forecast and forward price processes. We do this in Assumptions 3-4. These are natural assumptions in applications; they are satisfied by the demand forecast and forward price models used in §5.

Assumption 3 (Restricted forward price and demand forecast sets). In each stage $j \in \mathcal{J}$, the demand forecast D_j and the forward price F_j can take values in the bounded and finite sets $\mathcal{D}_j \subset \mathfrak{R}_+$ and $\mathcal{F}_j \subset \mathfrak{R}_+$, respectively.

Assumption 4 (Expanding reachable spot demand set). Define by $\widehat{\mathcal{D}}_j(D_j, F_j) \subseteq \mathcal{D}_{J+1}$ the set of spot demand values that can be reached in stage $J+1$ starting from $(D_j, F_j) \in \mathcal{D}_j \times \mathcal{F}_j$ in stage $j \in \mathcal{J}$; in particular, notice that $\widehat{\mathcal{D}}_{J+1}(d, s) \equiv \mathcal{D}_{J+1}$. Define by $\mathcal{N}_j(D_j, F_j) \subseteq \mathcal{D}_{j+1} \times \mathcal{F}_{j+1}$ the set of demand forecast and forward price values that can be reached with positive probability in stage $j+1$ starting from $(D_j, F_j) \in \mathcal{D}_j \times \mathcal{F}_j$ in stage $j \in \mathcal{J} \setminus \{J+1\}$. It holds that

$$\widehat{\mathcal{D}}_j(D_j, F_j) \equiv \bigcup_{(D_{j+1}, F_{j+1}) \in \mathcal{N}_j(D_j, F_j)} \widehat{\mathcal{D}}_{j+1}(D_{j+1}, F_{j+1}), \quad \forall j \in \mathcal{J} \setminus \{J+1\}, \quad (D_j, F_j) \in \mathcal{D}_j \times \mathcal{F}_j.$$

Proposition 4 characterizes the optimal value function and policy under this structure. We denote the maximum spot demand value by $\bar{d} := \max\{d : d \in \mathcal{D}_{J+1}\}$.

Proposition 4 (Optimal value function and policy with restriction). *Suppose that Assumptions 3-4 hold. In every stage $j \in \mathcal{J}$, the functions $V_j(x_j, D_j, F_j)$ and $U_j(x_j, D_j, F_j)$ are piecewise linear and continuous in $x_j \in \mathfrak{R}_+$ for each given $(F_j, D_j) \in \mathcal{D}_j \times \mathcal{F}_j$ with break points in set $\widehat{\mathcal{D}}_j(D_j, F_j)$. Moreover, in each stage $j \in \mathcal{J}$ the optimal basestock targets $\underline{b}_j(D_j, F_j)$ and $\bar{b}_j(D_j, F_j)$ belong to set $\widehat{\mathcal{D}}_j(D_j, F_j) \cup \{0, \bar{d}\}$.*

Proof. Consider stage $J+1$. Given $(d, s) \in \mathcal{D}_{J+1} \times \mathcal{F}_{J+1}$, the function $V_{J+1}(\cdot, d, s)$ changes slope in its first argument at d and the function $U_{J+1}(\cdot, d, s)$ is trivially piecewise linear and continuous in its first argument, being identically zero. Moreover, it holds that $\underline{b}_{J+1}(d, s) = \bar{b}_{J+1}(d, s) := d$. Thus, the claimed properties hold in stage $J+1$.

Consider stage J . Fix $(D_J, F_J) \in \mathcal{D}_J \times \mathcal{F}_J$. Given (D_J, F_J) , it follows from the stated assumptions that the function $U_J(x, D_J, F_J)$ is piecewise linear and continuous in its first argument and can change slope in this argument only at points in set $\widehat{\mathcal{D}}_J(D_J, F_J)$. The determination of $\underline{b}_J(D_J, F_J)$ involves solving the following optimization problem:

$$\max_{x_{J+1} \geq 0} U_J(x_{J+1}, D_J, F_J) - (1+B)F_J x_{J+1}. \quad (3)$$

The objective function of this problem can change slope in x_{J+1} only at points in set $\widehat{\mathcal{D}}_J(D_J, F_J)$.

Moreover, $x_{J+1} > \bar{d}$, say, $x_{J+1} = \bar{d} + \epsilon$ with $\epsilon > 0$, is better than setting $x_{J+1} = \bar{d}$ because

$$\begin{aligned}
U_J(\bar{d} + \epsilon, D_J, F_J) - (1 + B)F_J(\bar{d} + \epsilon) &= \mathbb{E}_J [(1 - A)s(\bar{d} + \epsilon - d)] - (1 + B)F_J(\bar{d} + \epsilon) \\
&= (1 - A)\mathbb{E}_J [s(\bar{d} - d)] - (1 + B)F_J\bar{d} + \{(1 - A)\mathbb{E}_J [s] \\
&\quad - (1 + B)F_J\}\epsilon \\
&= (1 - A)\mathbb{E}_J [s(\bar{d} - d)] - (1 + B)F_J\bar{d} - (A + B)F_J\epsilon \\
&< (1 - A)\mathbb{E}_J [s(\bar{d} - d)] - (1 + B)F_J\bar{d} \\
&= U_J(\bar{d}, D_J, F_J) - (1 + B)F_J\bar{d},
\end{aligned}$$

which is the objective function value of letting $x_{J+1} = \bar{d}$. Hence, the constraint $x_{J+1} \geq 0$ in (3) can be restricted to $x_{J+1} \in \widehat{\mathcal{D}}_J(D_J, F_J) \cup \{0, \bar{d}\}$ without loss of optimality.

Determining $\bar{b}_J(D_J, F_J)$ involves solving the optimization problem

$$\max_{x_{J+1} \geq 0} U_J(x_{J+1}, D_J, F_J) - (1 - B)F_Jx_{J+1}. \quad (4)$$

Setting $x_{J+1} > \bar{d}$ is dominated by letting $x_{J+1} = \bar{d}$ because for $\epsilon > 0$ it holds that

$$\begin{aligned}
U_J(\bar{d} + \epsilon, D_J, F_J) - (1 - B)F_J(\bar{d} + \epsilon) &= \mathbb{E}_J [(1 - A)s(\bar{d} + \epsilon - d)] - (1 - B)F_J(\bar{d} + \epsilon) \\
&= (1 - A)\mathbb{E}_J [s(\bar{d} - d)] - (1 - B)F_J\bar{d} - (A - B)F_J\epsilon \\
&< (1 - A)\mathbb{E}_J [s(\bar{d} - d)] - (1 - B)F_J\bar{d} \\
&= U_J(\bar{d}, D_J, F_J) - (1 - B)F_J\bar{d}.
\end{aligned}$$

Thus, the constraint $x_{J+1} \geq 0$ in (4) can be replaced with $x_{J+1} \in \widehat{\mathcal{D}}_J(D_J, F_J) \cup \{0, \bar{d}\}$ without loss of optimality.

It follows from Corollary 1 that $V_J(\cdot, D_J, F_J)$ is piecewise linear and continuous in its first argument and can change slope in this argument only at points in set $\widehat{\mathcal{D}}_J(D_J, F_J)$. Thus, the claimed properties also hold in stage J .

Make the induction hypothesis that these properties are also true in stages $j + 1$ through $J - 1$. Consider stage j . The induction hypothesis implies that the function $U_j(\cdot, D_j, F_j)$ can change slope in its first argument only at points in set $\widehat{\mathcal{D}}_j(D_j, F_j)$. Finding $\underline{b}_j(D_j, F_j)$ requires solving the following optimization problem:

$$\max_{x_{j+1} \geq 0} U_j(x_{j+1}, D_j, F_j) - (1 + B)F_jx_{j+1}. \quad (5)$$

Given an arbitrary $\epsilon > 0$, $x_{j+1} = \bar{d} + \epsilon$ is a better solution than \bar{d} because

$$\begin{aligned}
U_j(\bar{d} + \epsilon, D_j, F_j) - (1 + B)F_j(\bar{d} + \epsilon) &= \mathbb{E}_j[V_{j+1}(\bar{d} + \epsilon, D_{j+1}, F_{j+1})] - (1 + B)F_j(\bar{d} + \epsilon) \\
&= \mathbb{E}_j[(1 - B)F_{j+1}(\bar{d} + \epsilon) + \bar{K}_{j+1}(D_{j+1}, F_{j+1})] \\
&\quad - (1 + B)F_j(\bar{d} + \epsilon) \\
&\quad \text{(from Corollary 1 and } \underline{b}_{j+1}(D_{j+1}, F_{j+1}) \leq \bar{d}\text{)} \\
&= \mathbb{E}_j[(1 - B)F_{j+1}\bar{d} + \bar{K}_{j+1}(D_{j+1}, F_{j+1})] - (1 + B)F_j\bar{d} \\
&\quad - 2BF_j\epsilon \\
&< U_j(\bar{d}, D_j, F_j) - (1 + B)F_j\bar{d},
\end{aligned}$$

which is the objective function value associated with letting $x_{j+1} = \bar{d}$. Hence, the constraint $x_{j+1} \geq 0$ in (5) can be restricted to $x_{j+1} \in \widehat{\mathcal{D}}_j(D_j, F_j) \cup \{0, \bar{d}\}$ without loss of optimality.

Finding $\bar{b}_j(D_j, F_j)$ entails solving the following optimization problem:

$$\max_{x_{j+1} \geq 0} U_j(x_{j+1}, D_j, F_j) - (1 - B)F_j x_{j+1}. \quad (6)$$

Given an arbitrary $\epsilon > 0$, $x_{j+1} = \bar{d} + \epsilon$ is as good a solution as $x_{j+1} = \bar{d}$ because

$$\begin{aligned}
U_j(\bar{d} + \epsilon, D_j, F_j) - (1 - B)F_j(\bar{d} + \epsilon) &= \mathbb{E}_j[V_{j+1}(\bar{d} + \epsilon, D_{j+1}, F_{j+1})] - (1 - B)F_j(\bar{d} + \epsilon) \\
&= \mathbb{E}_j[(1 - B)F_{j+1}(\bar{d} + \epsilon) + \bar{K}_{j+1}(D_{j+1}, F_{j+1})] \\
&\quad - (1 - B)F_j(\bar{d} + \epsilon) \\
&= \mathbb{E}_j[(1 - B)F_{j+1}\bar{d} + \bar{K}_{j+1}(D_{j+1}, F_{j+1})] - (1 - B)F_j\bar{d} \\
&= U_j(\bar{d}, D_j, F_j) - (1 - B)F_j\bar{d},
\end{aligned}$$

which is the objective function value associated with letting $x_{j+1} = \bar{d}$. Hence, the constraint $x_{j+1} \geq 0$ in (6) can be restricted to $x_{j+1} \in \widehat{\mathcal{D}}_j(D_j, F_j) \cup \{0, \bar{d}\}$ without loss of optimality. Corollary 1 implies that the function $V_j(\cdot, D_j, F_j)$ is piecewise linear and continuous in its first argument and can change slope in this argument only at points in set $\widehat{\mathcal{D}}_j(D_j, F_j)$. Thus, the claimed properties hold in stage j . The principle of mathematical induction implies that they hold in every stage. \square

Proposition 4 is useful because it implies that in solving model (1)-(2) it is sufficient to compute the optimal value function for a finite set of feasible inventory levels, as opposed to an infinite number of such levels (those in \mathfrak{R}_+). Moreover, this can be done by limiting the search for the optimal basestock levels in each stage and for each demand forecast and forward price pair to a finite set of inventory levels. For example, it is sufficient to consider the unique inventory levels in

set $\{0\} \cup \mathcal{D}_{J+1}$ (recall that $\bar{d} \in \mathcal{D}_{J+1}$), although one can further refine the choice of which inventory levels to consider in each stage and for each demand forecast and forward price pair.

5. Numerical Study

In this section we report the results of a computational study aimed at assessing the effect of incorporating demand forecast and forward price updates in making commodity procurement choices. Before discussing these results in §5.4, we present the specific forward price and demand forecast models employed in this study in §5.1, the instances used as test beds in §5.2, and the basestock policies employed in this study together with their computation and evaluation in §5.3.

5.1 Demand Forecast and Forward Price Models

We employ versions of the celebrated futures price evolution model of Black [6] and the single product multiplicative version of the martingale model of (demand) forecast evolution of Heath and Jackson [29], which is essentially the model of Hausman [27]. Consistent with Black's model and Assumption 1, we assume that the forward price F_j evolves as a geometric Brownian motion with zero drift and positive volatility coefficient σ_F . The demand forecast also evolves as a geometric Brownian motion with volatility σ_D and zero drift, so that each demand forecast is unbiased; that is, the expected value of the spot demand is equal to the current demand forecast. When these two processes are correlated, they can be expressed as follows:

$$dD_j = \sigma_D D_j dZ_D \tag{7}$$

$$dF_j = \sigma_F F_j dZ_F \tag{8}$$

$$dZ_D dZ_F = \rho dt, \tag{9}$$

where dZ_D and dZ_F are correlated increments to standard Brownian motions with instantaneous correlation coefficient ρ , and dt is an infinitesimal time increment. These processes satisfy Assumption 2. We use discrete time and space versions of model (7)-(9) in our numerical study. Specifically, we implement the joint evolution of the price and demand forecast processes using a tridimensional tree using the code available in Haugh [26, §3.3]. Thus, Assumptions 3-4 hold in our implementation, and so does Proposition 4.

5.2 Instances

We use natural gas data augmented with data available in the literature. We consider stage $J + 1$ to correspond to March 2010. We use \$5.591/MMBtu as the initial forward price; this is the closing

price on August 14, 2009, of the NYMEX natural gas futures contract for delivery in March 2010.

In our study, the firm is a hypothetical natural gas distributor operating in the U.S. We take the firm's March 2010 demand forecast as of the beginning of the time horizon to be 14,403,838MMBtu, which corresponds to the March 2002 demand faced by the utility considered in the study of Muthuraman et al. [48, Table 1, p. 1143]. Here, we make the convenient assumption that this figure is indicative of the demand forecast of this firm for March 2010.

We let 6 months be the base case for the length of the time horizon, which is equivalent to letting the initial stage being the beginning of September 2009. Denoting by T_j the length of time in between stages j and $J + 1$, this approximately correspond to setting T_1 equal to 180/365. We also consider two additional values for the time horizon parameter, 2 and 4 months, which correspond to values of 60/365 and 120/365 for T_1 .

We set a base value for the demand forecast volatility to make the CV of spot demand corresponding to the 6 month horizon consistent with the CV of spot demand used by Seifert et al. [55], which is 0.25. Given our assumptions, conditional on the demand forecast being D_j in stage j , the distribution of the spot demand d is lognormal with mean D_j and standard deviation $D_j \sqrt{\exp(\sigma_D^2 T_j) - 1}$, so that the CV of spot demand given the information available in stage j is equal to $\sqrt{\exp(\sigma_D^2 T_j) - 1}$. Thus, for given CV and T_j values, the corresponding volatility parameter σ_D is equal to $\sqrt{\ln(1 + CV^2)/T_j}$. For $T_1 = 180/365$ and $CV = 0.25$, we obtain a value for σ_D equal to 0.3506, which we round to 0.35. We also consider values for σ_D equal to 0.21 and 0.48, which roughly correspond to spot demand CV values equal to 0.15 and 0.35 when $T_1 = 180/365$.

For the forward price volatility, we consider a base value of 0.60, which is roughly consistent with the implied volatility for the 6 month maturity of the Summer case reported by Lai et al. [43, Table 11, p. OA-3, Online Appendix]. We also consider the two following additional values for this parameter: 0.4 and 0.8.

Similar to the demand forecast volatility, we set a base value for the instantaneous correlation coefficient ρ by making the correlation between the spot price and the spot demand corresponding to the 6 month horizon consistent with the analogous correlation value used by Seifert et al. [55], which is 0.2. Specifically, given our assumptions, the correlation between the spot price and the spot demand given the information available in stage j is $[\exp(\sigma_D \sigma_F \rho T_j) - 1]/(C_D C_F)$, with $C_D := \sqrt{\exp(\sigma_D^2 T_j) - 1}$ and $C_F := \sqrt{\exp(\sigma_F^2 T_j) - 1}$. Thus, for given values of this correlation, denoted by CORR, and T_j , the corresponding instantaneous correlation coefficient ρ is equal to $\ln(1 + C_D C_F \text{CORR})/(\sigma_D \sigma_F T_j)$. For $T_1 = 180/365$, $\sigma_D = 0.35$, and $\sigma_F = 0.60$ we obtain $\rho = 0.2101$, which we round to 0.21. We also consider values for ρ equal to 0.42 and 0.62, which roughly

Table 1: Policies Considered in the Numerical Study.

Policy	Optimization	Update	
		Demand Forecast	Forward Price
D1			
D2		✓	
O1	✓		
O2	✓		✓
O3	✓	✓	✓

correspond to CORR values equal to 0.40 and 0.60, based on the argument made by Seifert et al. [55] that the correlation between a commodity price and the demand for this commodity faced by a firm should be positive.

We generate our test instances by considering all the possible combinations of the stated parameters. In terms of the transaction costs, Muthuraman et al. [48] use a brokerage fee of \$0.30/MMBtu and a futures price equal to \$9.00/MMBtu, which is consistent with setting the coefficient B equal to $1/30$. The analysis of Williams [66, p. 1014] suggests that the spot price transaction cost is roughly 3 times as large as the forward price transaction cost. Thus, we set A equal to $1/10$. We use these values for these parameters in all the considered instances.

5.3 Basestock Policies

In our study we compare the performance of different basestock policies, labeled D1, D2, O1, O2, and O3, which differ in terms of whether they are static or dynamic and the type of information they use to compute their basestock levels (see Table 1). The static policies follow:

- *D1 policy.* This policy makes an initial procurement in the forward market equal to the demand forecast in the first stage and state, does nothing until the last stage, when it trades in the spot market to make up any differences between the amount of commodity on hand and the realized spot demand. Thus, this is a basestock policy whose basestock levels in the initial stage and state are equal to the demand forecast, and the buy-up-to and sell-down-to levels are equal to zero and \bar{d} , respectively, in each remaining stage and state, excluding the last stage where they are equal to the realized spot demand in each state. This policy reflects current practice at the energy reseller whose practices motivated this research.
- *O1 policy.* This policy is essentially a newsvendor model that optimizes the amount of supply to purchase in the forward market in the initial stage in the face of spot price and demand uncertainty, makes no further trades in the forward market, and makes up any differences

between the procured supply and the spot demand in the spot market in the final stage. Thus, this policy is analogous to the D1 policy except that its basestock levels in the initial stage and state are determined by the newsvendor optimal solution, and hence depend on both the forward price and the demand forecast.

We use these static policies as benchmarks for the following dynamic policies:

- *D2 policy.* This policy is a dynamic version of the D1 policy that in every stage and state adjusts the total amount of procured supply by making it equal to the current demand forecast; hence, its basestock levels are equal to the demand forecast in every stage and state. This policy is conceptually consistent with the policy proposed by Levary and Dean [45], because it reactively accounts for the newly available demand forecasts updates, although their model uses optimization.
- *O2 policy.* This policy extends the O1 policy by trading in the forward market in every stage and state, by taking into account the forward price updates but not the demand forecast updates. Thus, the basestock levels of this policy only depend on the forward price in each stage and state, except that in the first stage they depend on both price and demand information, and in the last stage they depend only on the realized spot demand. This policy is loosely representative of procurement models discussed in the real options literature, which only take dynamic price information into account; however, these models may also take capacity constraints into account, which are not considered in this study.
- *O3 policy.* This is the optimal policy to model (1)-(2), analyzed in §4. This policy extends the O2 policy by taking into account the demand forecast update in every stage and state. It can be conceptually interpreted as extending real options models by adding to them another source of dynamically evolving information, the demand forecast.

We compute and evaluate all of these policies based on the tridimensional tree discussed at the end of §5.2. Specifically, the code we use to generate this tridimensional tree requires as input a time step, as well as the length of the time horizon, the starting values for the demand forecast and forward price, their respective volatilities, and the instantaneous correlation coefficient. We let the time step be $10/365$, which corresponds to a 10 day increment normalized by the number of days in one year; the choice of the other parameters is described in §5.2. For each combination of these parameters, we use each generated tridimensional tree as a starting point to compute the various

policies and their values in the initial stage and state (we assume that the initial inventory level is zero).

In particular, for the O3 policy we use each such tree to define the stages of each MDP and its price and demand information states in each stage. In addition, we consider a number of inventory levels in each stage as discussed after Proposition 4. We optimally solve each such MDP by backward dynamic programming, by leveraging the structure of the optimal policy established in Proposition 3 to speed up the computation of the basestock levels. This also provides the value of this policy in the initial stage and state.

We compute and evaluate the O2 policy in an analogous manner, except that its state space in every stage, excluding the last one, does not include the updated demand forecast information. For brevity we do not provide the formulation of the resulting MDP. We only point out that in transitioning from stage J to stage $J + 1$, when the spot demand information becomes known, we compute the relevant expectation conditional on the realized forward price in stage J and the initial value of the demand forecast in stage 1.

We compute and evaluate the O1 policy by using the joint spot price and demand distribution available in the initial stage. This distribution corresponds to all the spot price and demand levels that can be reached in stage $J + 1$ on the tridimensional tree and their associated probabilities, which we compute by means of a straightforward forward recursion on this tree. The initial procurement choice is limited to a value in the demand support of this distribution or zero.

We also compute and evaluate the D1 policy in this manner, that is, by setting the initial forward procurement choice equal to the value in the demand support of this distribution or zero that is closest to the initial demand forecast. We compute and evaluate the D2 policy by embedding and repeating this process in a forward recursion on the tridimensional tree. In this forward recursion we update the joint distribution of the spot price and demand in every stage and state, and keep track of both the probability to reach a given state in a given stage and the value of the procurement policy associated with reaching this state and stage, that is, the probability-to-come and the value-to-come to a stage and state. This generates a distribution for the random variable value-to-come to a state in the last stage (the relevant probabilities here are the probabilities-to-come to each state in the last stage). The value of the D2 policy is the expected value of this random variable.

5.4 Results

In discussing our computational results, we focus on assessing the effects of incorporating information updates in devising a commodity procurement policy. We do this by separately measuring

the effects of considering demand forecast updates but not forward price updates, forward price updates but not demand forecast updates, and both types of updates. We measure these effects twice, once taking the D1 policy as benchmark and once taking the O1 policy as benchmark, because this allows us to quantify these effects relative to the cases when optimization is and is not used to compute a static procurement policy, respectively. We quantify these effects by restricting ourselves to the basestock type policies discussed in §5.3. Specifically, we measure the demand forecast update effect, the forward price update effect, and the combined demand forecast and forward price effect by comparing the performances of the D2, O2, and O3 policies, respectively, to that of each of the two benchmark policies. Moreover, we discuss the sensitivities of these effects to parameters of interest and the potential managerial relevance of these effects.

Information update effects. Tables 2, 3, and 4 report the magnitudes of these effects for values of the time horizon equal to two, four, and six months, respectively, for each combination of the considered demand and price volatilities and instantaneous correlation coefficient. (Tables 5, 6, and 7 in Appendix B report the raw data used to compute these figures.) Figure 1 graphically illustrates the minimum and maximum values of the effects, which are also summarized at the bottom of Tables 2, 3, and 4. It is important to emphasize that the figures displayed in Tables 2, 3, and 4 do not depend on the values chosen for the initial forward price and demand forecast (this can be easily shown). In this sense, the percent improvements displayed in these tables are generic.

Consider first the case when the D1 policy is the benchmark. The demand forecast effect is negative for all the considered time horizons. The increase in the cost of the procurement policy varies between 0.04-0.47%, 0.27-1.35%, and 0.54-2.32% in the two, four, and six month cases, respectively. Although it may appear surprising that taking more information into account can be detrimental, this occurs because adjusting the procured supply to simply match the revised demand forecast in each stage and state is costly on average, as it incurs the transaction costs and the forward price is a martingale (constant on average). The forward price update effect is small but positive, improving the value of the procurement policy by 0.04-0.38%, 0.10-0.49%, and 0.13-0.59%, respectively, in the three time horizon cases. In contrast, the combined forward price and demand forecast update effect is substantially larger, with associated improvements equal to 0.37-0.98%, 0.63-1.48%, and 0.80-1.92%, for the three considered time horizons.

When the O1 policy is taken as benchmark, relative to the case when the benchmark is the D1 policy, the demand forecast update effect slightly increases (in absolute value, as it remains negative), and the forward price and combined demand forecast and forward price update effects slightly decrease, due to the fact that O1 is a better policy than D1. Specifically, the former effect

Table 2: The Effects of Information Updating: Two Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	Base: D1			Base: O1		
	Demand	Forward		Demand	Forward	
	Forecast (D1)	Price (O2)	Both (O3)	Forecast (D1)	Price (O2)	Both (O3)
(0.21, 0.40, 0.21)	-0.22	0.04	0.39	-0.25	0.01	0.36
(0.21, 0.40, 0.42)	-0.15	0.09	0.40	-0.21	0.04	0.34
(0.21, 0.40, 0.62)	-0.05	0.14	0.38	-0.07	0.11	0.35
(0.21, 0.60, 0.21)	-0.22	0.04	0.39	-0.25	0.01	0.36
(0.21, 0.60, 0.42)	-0.15	0.09	0.39	-0.20	0.04	0.35
(0.21, 0.60, 0.62)	-0.04	0.13	0.37	-0.06	0.12	0.36
(0.21, 0.80, 0.21)	-0.21	0.05	0.40	-0.25	0.02	0.37
(0.21, 0.80, 0.42)	-0.15	0.09	0.40	-0.19	0.05	0.36
(0.21, 0.80, 0.62)	-0.04	0.13	0.38	-0.05	0.12	0.37
(0.35, 0.40, 0.21)	-0.36	0.07	0.64	-0.42	0.01	0.58
(0.35, 0.40, 0.42)	-0.25	0.15	0.65	-0.35	0.06	0.55
(0.35, 0.40, 0.62)	-0.07	0.29	0.68	-0.17	0.18	0.58
(0.35, 0.60, 0.21)	-0.36	0.07	0.64	-0.41	0.01	0.58
(0.35, 0.60, 0.42)	-0.25	0.14	0.64	-0.33	0.06	0.56
(0.35, 0.60, 0.62)	-0.07	0.27	0.66	-0.15	0.19	0.58
(0.35, 0.80, 0.21)	-0.35	0.07	0.64	-0.40	0.02	0.59
(0.35, 0.80, 0.42)	-0.25	0.14	0.64	-0.31	0.07	0.57
(0.35, 0.80, 0.62)	-0.06	0.26	0.65	-0.13	0.19	0.59
(0.48, 0.40, 0.21)	-0.47	0.22	0.98	-0.68	0.01	0.77
(0.48, 0.40, 0.42)	-0.35	0.20	0.87	-0.47	0.08	0.74
(0.48, 0.40, 0.62)	-0.09	0.38	0.91	-0.23	0.24	0.77
(0.48, 0.60, 0.21)	-0.47	0.21	0.97	-0.67	0.01	0.78
(0.48, 0.60, 0.42)	-0.34	0.19	0.85	-0.45	0.08	0.75
(0.48, 0.60, 0.62)	-0.09	0.36	0.89	-0.20	0.25	0.78
(0.48, 0.80, 0.21)	-0.47	0.20	0.96	-0.65	0.02	0.79
(0.48, 0.80, 0.42)	-0.34	0.18	0.85	-0.43	0.09	0.76
(0.48, 0.80, 0.62)	-0.08	0.34	0.87	-0.18	0.24	0.78
Minimum	-0.47	0.04	0.37	-0.68	0.01	0.34
Maximum	-0.04	0.38	0.98	-0.05	0.25	0.79

Table 3: The Effects of Information Updating: Four Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	Base: D1			Base: O1		
	Demand	Forward		Demand	Forward	
	Forecast (D1)	Price (O2)	Both (O3)	Forecast (D1)	Price (O2)	Both (O3)
(0.21, 0.40, 0.21)	-0.59	0.10	0.66	-0.67	0.02	0.58
(0.21, 0.40, 0.42)	-0.48	0.15	0.65	-0.56	0.07	0.58
(0.21, 0.40, 0.62)	-0.29	0.22	0.64	-0.34	0.16	0.58
(0.21, 0.60, 0.21)	-0.59	0.10	0.66	-0.66	0.03	0.59
(0.21, 0.60, 0.42)	-0.48	0.14	0.64	-0.54	0.08	0.58
(0.21, 0.60, 0.62)	-0.28	0.21	0.63	-0.31	0.17	0.59
(0.21, 0.80, 0.21)	-0.59	0.11	0.67	-0.65	0.05	0.61
(0.21, 0.80, 0.42)	-0.47	0.14	0.65	-0.52	0.10	0.60
(0.21, 0.80, 0.62)	-0.27	0.21	0.63	-0.29	0.19	0.61
(0.35, 0.40, 0.21)	-0.98	0.19	1.09	-1.14	0.04	0.94
(0.35, 0.40, 0.42)	-0.79	0.26	1.08	-0.94	0.11	0.93
(0.35, 0.40, 0.62)	-0.48	0.36	1.04	-0.58	0.26	0.94
(0.35, 0.60, 0.21)	-0.98	0.18	1.08	-1.12	0.05	0.95
(0.35, 0.60, 0.42)	-0.78	0.24	1.06	-0.91	0.12	0.93
(0.35, 0.60, 0.62)	-0.46	0.33	1.02	-0.52	0.27	0.95
(0.35, 0.80, 0.21)	-0.97	0.18	1.09	-1.10	0.06	0.96
(0.35, 0.80, 0.42)	-0.77	0.23	1.05	-0.87	0.13	0.95
(0.35, 0.80, 0.62)	-0.43	0.32	1.01	-0.48	0.28	0.97
(0.48, 0.40, 0.21)	-1.35	0.29	1.48	-1.59	0.05	1.25
(0.48, 0.40, 0.42)	-1.08	0.34	1.43	-1.28	0.15	1.23
(0.48, 0.40, 0.62)	-0.65	0.49	1.40	-0.80	0.34	1.25
(0.48, 0.60, 0.21)	-1.34	0.27	1.46	-1.56	0.06	1.25
(0.48, 0.60, 0.42)	-1.07	0.31	1.40	-1.23	0.15	1.24
(0.48, 0.60, 0.62)	-0.62	0.45	1.36	-0.73	0.34	1.26
(0.48, 0.80, 0.21)	-1.34	0.27	1.46	-1.53	0.07	1.27
(0.48, 0.80, 0.42)	-1.05	0.29	1.38	-1.18	0.17	1.25
(0.48, 0.80, 0.62)	-0.59	0.42	1.34	-0.66	0.36	1.27
Minimum	-1.35	0.10	0.63	-1.59	0.02	0.58
Maximum	-0.27	0.49	1.48	-0.29	0.36	1.27

Table 4: The Effects of Information Updating: Six Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	Base: D1			Base: O1		
	Demand	Forward		Demand	Forward	
	Forecast (D1)	Price (O2)	Both (O3)	Forecast (D1)	Price (O2)	Both (O3)
(0.21, 0.40, 0.21)	-1.03	0.13	0.86	-1.14	0.03	0.76
(0.21, 0.40, 0.42)	-0.87	0.17	0.83	-0.95	0.09	0.75
(0.21, 0.40, 0.62)	-0.59	0.25	0.82	-0.65	0.19	0.76
(0.21, 0.60, 0.21)	-1.03	0.13	0.86	-1.12	0.04	0.77
(0.21, 0.60, 0.42)	-0.86	0.15	0.82	-0.92	0.10	0.76
(0.21, 0.60, 0.62)	-0.57	0.24	0.80	-0.60	0.21	0.77
(0.21, 0.80, 0.21)	-1.02	0.15	0.87	-1.10	0.07	0.80
(0.21, 0.80, 0.42)	-0.84	0.16	0.83	-0.89	0.12	0.79
(0.21, 0.80, 0.62)	-0.54	0.25	0.81	-0.56	0.23	0.80
(0.35, 0.40, 0.21)	-1.71	0.26	1.43	-1.93	0.05	1.21
(0.35, 0.40, 0.42)	-1.44	0.30	1.37	-1.60	0.14	1.21
(0.35, 0.40, 0.62)	-0.98	0.42	1.33	-1.09	0.30	1.21
(0.35, 0.60, 0.21)	-1.70	0.25	1.41	-1.89	0.06	1.22
(0.35, 0.60, 0.42)	-1.42	0.27	1.34	-1.54	0.14	1.21
(0.35, 0.60, 0.62)	-0.94	0.38	1.29	-1.01	0.31	1.23
(0.35, 0.80, 0.21)	-1.69	0.25	1.42	-1.86	0.08	1.25
(0.35, 0.80, 0.42)	-1.39	0.26	1.33	-1.49	0.17	1.24
(0.35, 0.80, 0.62)	-0.89	0.37	1.29	-0.93	0.34	1.25
(0.48, 0.40, 0.21)	-2.32	0.38	1.92	-2.66	0.06	1.59
(0.48, 0.40, 0.42)	-1.96	0.43	1.83	-2.21	0.17	1.59
(0.48, 0.40, 0.62)	-1.33	0.59	1.80	-1.53	0.39	1.60
(0.48, 0.60, 0.21)	-2.31	0.36	1.89	-2.61	0.07	1.60
(0.48, 0.60, 0.42)	-1.93	0.37	1.79	-2.13	0.18	1.59
(0.48, 0.60, 0.62)	-1.28	0.53	1.74	-1.41	0.40	1.61
(0.48, 0.80, 0.21)	-2.30	0.35	1.89	-2.57	0.09	1.63
(0.48, 0.80, 0.42)	-1.89	0.35	1.76	-2.04	0.20	1.62
(0.48, 0.80, 0.62)	-1.21	0.49	1.71	-1.29	0.42	1.63
Minimum	-2.32	0.13	0.80	-2.66	0.03	0.75
Maximum	-0.54	0.59	1.92	-0.56	0.42	1.63




Figure 1: Summary of the numerical results pertaining to the information update effects.

ranges (in absolute value) between 0.05-0.68%, 0.29-1.59%, and 0.56-2.66% in the three considered time horizons, and the latter effects between 0.01-0.25%, 0.02-0.36%, 0.03-0.42% and 0.34-0.79%, 0.58-1.27%, and 0.75-1.63% in the same time horizons.

Benefit of dynamic optimization. These results can be interpreted in terms of the benefit of dynamic optimization. There is some benefit from dynamic, as opposed to static, optimization even if only forward price updates are considered. Put differently, static optimization does not provide all the benefit of dynamic optimization when the latter excludes demand forecast updates; it can be verified that the O1 policy improves the D1 policy by 0.01-0.21%, 0.02-0.24%, and 0.01-0.33% in the two, four, and six month time horizons, respectively. But dynamic optimization is substantially more useful than static optimization when one also considers demand forecast updates; that is, the O3 policy improves the O1 policy substantially more than the O2 policy improves the O1 policy.

Sensitivity analysis. We observe that increasing the time horizon reinforces all the considered effects, no matter what the benchmark policy. The same occurs when the demand forecast volatility increases. This can be explained by noting that the CV of spot demand increases in the demand volatility and the length of the horizon, which confirms a familiar insight in inventory management. In almost all the cases, the forward price update effect increases with increasing instantaneous correlation, as with larger values of this parameter a forward price update carries more information about the demand forecast process. The demand forecast update effect decreases in absolute value with increasing instantaneous correlation. This occurs because the costs (absolute values) of the D1 and the D2 policies increase as the instantaneous correlation increases, but the gap between the costs of the D1 and the D2 policies reduces. For a similar reason, the combined effect decreases in the instantaneous correlation when the benchmark policy is D1; but it is rather insensitive to changes in this parameter when the benchmark policy is O1. The considered effects are largely unaffected by changes in the forward price volatility.

Summary. The main insights from this analysis can be summarized as follows. With transaction costs, when to update a firm's supply position in the forward market is nontrivial. Price information and optimization are required to take advantage of information updates (with optimization one can always discard new information if this is not beneficial). Intriguingly, although ignoring price information is detrimental when considering demand forecast updates, the effect of considering forward price updates alone is small, but the combined demand forecast and forward price update effect is substantially larger. This suggests that there is a complementary relationship between these two types of information updates, in the sense that one would not want to consider one type of information update without considering the other. Moreover, dynamic optimization

is substantially more useful than static optimization in the presence of demand forecast updates. Increasing demand variability reinforces these conclusions.

Potential managerial relevance. Even if larger than the other effects, the combined demand forecast and forward price update effect is not large when expressed as a percentage. Indeed this is expected given that the forward price and the demand forecast are martingales and the magnitude of the considered transaction costs. To assess the potential managerial relevance of this combined effect, it is useful to consider their associated absolute improvement figures.

Taking the D1 policy as benchmark, the absolute improvements of the combined effect vary between \$0.3-0.8M, \$0.5-1.3M, and \$0.7-1.7M in the two, four, and six month time horizons; with the O1 policy as benchmark, these improvement ranges are \$0.3-0.7M, \$0.5-1.2M, and \$0.6-1.6M. Given that these improvements pertain to a single month, they seem to have significant potential managerial relevance. In particular, they should be relevant to unregulated energy resellers that operate with small margins, the typical case in practice, because these are margin improvements.

They should also be relevant to regulated local distribution companies that operate under benchmark based regulation (see, e.g., Muthuraman et al. [48]). In this case the local distribution company's margin amounts to a bonus proportional to the difference between the cost of a benchmark procurement policy set by the regulator and the cost of the procurement policy employed by the utility; the constant of proportionality is the bonus rate. Given that the benchmark policy set by the regulator is unlikely to depend on the demand forecast update process of a local distribution company, a fraction, equal to the bonus rate, of any cost reduction obtained by incorporating this information in the procurement policy of such a firm would amount to a margin improvement for this company, with the rest of the improvement benefiting consumers.

6. Conclusions

In this paper we consider the problem of procuring a commodity to satisfy an uncertain requirement (demand) at a given future date, when trading can be done both in forward and spot markets that feature differential transaction costs, and the forward price and the demand forecasts evolve over time as correlated stochastic processes. We formulate this problem as an MDP, establish that its optimal procurement policy is of the basestock type, and discuss its computation in cases of practical relevance.

In a numerical study based on natural gas data, we find that incorporating demand forecast updates in a firm's procurement policy is nontrivial. We find it intriguing that the combined effect of using both forward price and demand forecast updates in a firm's procurement policy is

substantially larger than the effect of incorporating only forward price updates in this policy; this suggest a complementarity relationship between the two types of updates. Increased demand variability reinforces these effects, which are instead largely insensitive to increased price variability. The correlation between price and demand does not substantially impact the combined effect, but the individual effects are more sensitive to changes in this parameter. Fine tuning a firm’s procurement policy by accounting for the combined effect has potential managerial relevance, offering improvements between \$0.3-1.7M per month, depending on the benchmark policy and the specific parameters used. These insights and results have particular significance for energy resellers and local distribution companies, but retain relevance in other industrial and commercial contexts.

Our work is limited to a single future date when demand is to be satisfied. It would be of interest to consider multiple such dates, in which case commodity storage (inventory) considerations would become important. In this case it would also be important to model the joint evolution of the forward curve for the commodity and the constellation of demand forecasts. These aspects are likely to pose significant challenges in terms of model tractability, especially from a computational standpoint. In this case, developing effective approximate dynamic programming heuristics and tight performance bounds appear to be fruitful areas for additional research.

Our model does not feature capacity constraints. These may pertain to limits in the amount of trading that can be performed in the forward market during a given time period; the relevant limitation here would likely arise from working capital availability rather than limits in the number of trades that one is able to perform. Our model could be extended in this direction, possibly at the expense of including additional constraints and/or state variables in the MDP formulation. This is challenging because proper modeling of these features should be done in the context of specific applications, to reflect the contextual richness of the application domain. The resulting MDP would resemble a generalized swing option model that includes both price and demand uncertainty. It would be interesting to assess the relative magnitudes of the information update effects studied here. These generalizations could be explored in future research.

A. On the Application of Risk Neutral Valuation

If financial markets were arbitrage free, complete, and perfect, then one could apply risk neutral valuation theory (Cox et al. [12], Duffie [15], Hull [32, Chapter 25]) and specify model (1)-(2) using risk neutral versions of the forward price and the firm’s demand forecast stochastic processes; in particular when they are modeled as diffusion processes, as we do in §5.

Lack of arbitrage in financial markets is a necessary condition for application of this approach.

Financial market completion would occur if the three following conditions were satisfied: the risk free rate were deterministic and constant, there existed a futures market for the commodity, as well as financial claims whose payoffs are contingent on realizations of the firm's demand forecast process. The first condition is standard and implies equivalence between futures and forward prices (Cox et al. [11]). The second condition is realistic for many commodities of interest, in particular energy sources such as natural gas. The third condition would be satisfied, for example, if the firm's demand forecast process were perfectly correlated with the evolution of the temperature at a particular location (this is not too far from reality for residential and commercial demand for natural gas), and there existed a market for weather derivatives written on this temperature (such derivatives are listed on the Chicago Mercantile Exchange; see also Dischel [13] and Hull [32, Chapter 23]). If the third condition failed to apply, one could still use risk neutral valuation as an approximation, as discussed by Smith [59, p. 95]; in this case, assuming that changes in the firm's demand forecasts (opportune modeled, for example as we do in §5) were uncorrelated with changes in the price of the market portfolio would require no risk adjustment to the firm's demand forecast process; this is fairly realistic for energy demand as discussed by Hull [32, Chapters 23].

Financial market perfection requires assuming that trading in financial markets, e.g., the stated commodity futures and weather derivatives markets, does not incur transaction costs, that is, the firm incurs such costs only when trading in physical markets, such as the spot and forward markets for the commodity. If instead trading in financial markets commands transaction costs, then one could use risk neutral valuation heuristically, by specifying opportune stochastic processes for the relevant derivative prices. For example, given that the risk neutral versions of these processes are martingales when there are no transaction costs, versions of these processes perturbed to account for transaction costs would be arbitrage free in the framework introduced by Jouini and Kallal [34] (see in particular their Theorem 3.2 and the discussion immediately following the statement of this theorem). Applying risk neutral valuation using the martingale versions of these processes to determine a procurement policy would be consistent with the no arbitrage framework of these authors, and would value this policy at one of the possible values that do not admit arbitrage (see part (iii) of the stated theorem).

B. Raw Data

Tables 5-7 present the raw data used to perform the analysis discussed in §5. The displayed values are costs (absolute values of the policy values) because the value of each considered policy in the initial stage and state is negative.

Table 5: Policy Costs in the Initial Stage and State: Two Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	D1	D2	O1	O2	O3
(0.21, 0.40, 0.21)	83988853.90	84170611.50	83956737.40	83951543.42	83657397.01
(0.21, 0.40, 0.42)	84240146.13	84367528.40	84192085.27	84160991.67	83904469.43
(0.21, 0.40, 0.62)	84449358.49	84487507.68	84428990.90	84332596.80	84129917.28
(0.21, 0.60, 0.21)	84106454.40	84287586.71	84077427.20	84069631.42	83774501.08
(0.21, 0.60, 0.42)	84475434.37	84602119.13	84434372.39	84400544.14	84142114.15
(0.21, 0.60, 0.62)	84799990.61	84834681.55	84787434.68	84688276.27	84482967.22
(0.21, 0.80, 0.21)	84223547.47	84403872.70	84197217.89	84182401.76	83886343.95
(0.21, 0.80, 0.42)	84710146.30	84835655.42	84675654.01	84634394.89	84374109.85
(0.21, 0.80, 0.62)	85150580.96	85180713.54	85141216.51	85038160.53	84830263.97
(0.35, 0.40, 0.21)	84492827.98	84795010.63	84440059.95	84432135.25	83951641.80
(0.35, 0.40, 0.42)	84914245.27	85129255.68	84835126.92	84785115.97	84364474.23
(0.35, 0.40, 0.62)	85321259.36	85380367.63	85232522.53	85075883.71	84741805.65
(0.35, 0.60, 0.21)	84690070.21	84991128.30	84642372.15	84631912.70	84149347.12
(0.35, 0.60, 0.42)	85309654.44	85523315.21	85242024.65	85189180.15	84764584.91
(0.35, 0.60, 0.62)	85905920.87	85964155.47	85835531.40	85675723.37	85336325.49
(0.35, 0.80, 0.21)	84886592.60	85186232.05	84843322.76	84826011.92	84341468.83
(0.35, 0.80, 0.42)	85704689.65	85916230.96	85647847.12	85587401.76	85158935.90
(0.35, 0.80, 0.62)	86491882.11	86547384.36	86432242.54	86270322.92	85925660.24
(0.48, 0.40, 0.21)	85056109.85	85454804.80	84880009.29	84869769.05	84224013.49
(0.48, 0.40, 0.42)	85532434.90	85830251.20	85425471.27	85359006.04	84791589.03
(0.48, 0.40, 0.62)	86094267.54	86175521.59	85973636.28	85764028.85	85311073.95
(0.48, 0.60, 0.21)	85322770.12	85725324.05	85158914.60	85146203.60	84497081.43
(0.48, 0.60, 0.42)	86078830.32	86374489.34	85987357.11	85917919.30	85344132.92
(0.48, 0.60, 0.62)	86903426.85	86982215.95	86807608.50	86594219.72	86132808.80
(0.48, 0.80, 0.21)	85588724.59	85994666.33	85436130.64	85416737.36	84764377.82
(0.48, 0.80, 0.42)	86625468.87	86917935.45	86548544.47	86471320.97	85891269.43
(0.48, 0.80, 0.62)	87716076.63	87789863.02	87634867.41	87420706.87	86950890.86

Table 6: Policy Costs in the Initial Stage and State: Four Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	D1	D2	O1	O2	O3
(0.21, 0.40, 0.21)	84458443.17	84958549.78	84392198.14	84371197.84	83899053.61
(0.21, 0.40, 0.42)	84944969.17	85353282.40	84880275.60	84819495.65	84390645.29
(0.21, 0.40, 0.62)	85396820.57	85643038.50	85349602.05	85210232.09	84850718.57
(0.21, 0.60, 0.21)	84694298.61	85194153.49	84635297.46	84608087.48	84134204.86
(0.21, 0.60, 0.42)	85420583.81	85827291.37	85368039.94	85302267.45	84870111.79
(0.21, 0.60, 0.62)	86106555.17	86346363.82	86076168.27	85929603.34	85565744.93
(0.21, 0.80, 0.21)	84929393.85	85428639.13	84875818.64	84833721.98	84358204.97
(0.21, 0.80, 0.42)	85896770.67	86300412.98	85855754.64	85773961.50	85338591.27
(0.21, 0.80, 0.62)	86820096.36	87050182.37	86801142.97	86639274.58	86271133.78
(0.35, 0.40, 0.21)	85285758.33	86123791.24	85156865.22	85124375.08	84356329.16
(0.35, 0.40, 0.42)	86107173.69	86789509.05	85978924.73	85882327.71	85180693.10
(0.35, 0.40, 0.62)	86858993.93	87272776.09	86771622.64	86547252.68	85954765.32
(0.35, 0.60, 0.21)	85681755.79	86520659.74	85565855.17	85526995.53	84754629.45
(0.35, 0.60, 0.42)	86909196.42	87591118.42	86803276.02	86701957.85	85992170.51
(0.35, 0.60, 0.62)	88062328.52	88465941.74	88004339.75	87770843.74	87167750.01
(0.35, 0.80, 0.21)	86077163.49	86916341.58	85971241.08	85918219.35	85141705.64
(0.35, 0.80, 0.42)	87714961.93	88394007.34	87630255.53	87512653.42	86794838.19
(0.35, 0.80, 0.62)	89278181.98	89666063.63	89240776.81	88991291.95	88377630.05
(0.48, 0.40, 0.21)	86054215.86	87212365.28	85849751.81	85807832.52	84780560.97
(0.48, 0.40, 0.42)	87160469.95	88104997.56	86988063.46	86861369.38	85917687.54
(0.48, 0.40, 0.62)	88222518.70	88795862.77	88088366.45	87791182.67	86988539.50
(0.48, 0.60, 0.21)	86600603.46	87762071.38	86415408.82	86366994.94	85332085.90
(0.48, 0.60, 0.42)	88275458.10	89219038.71	88132967.34	88001572.85	87043484.10
(0.48, 0.60, 0.62)	89898395.57	90459151.39	89806472.27	89497693.77	88676328.56
(0.48, 0.80, 0.21)	87147099.84	88311024.54	86977037.76	86915144.30	85872826.31
(0.48, 0.80, 0.42)	89399088.79	90338440.29	89285021.77	89136979.41	88164598.58
(0.48, 0.80, 0.62)	91599552.16	92140082.84	91538102.03	91211975.24	90371813.86

Table 7: Policy Costs in the Initial Stage and State: Six Month Time Horizon.

$(\sigma_D, \sigma_F, \rho)$	D1	D2	O1	O2	O3
(0.21, 0.40, 0.21)	84871249.24	85748018.50	84785155.40	84759512.07	84141408.27
(0.21, 0.40, 0.42)	85590768.73	86335857.25	85524804.49	85449290.50	84879638.25
(0.21, 0.40, 0.62)	86282103.58	86793648.76	86232583.39	86064542.77	85577695.48
(0.21, 0.60, 0.21)	85226711.14	86103834.46	85150965.65	85116500.61	84495453.10
(0.21, 0.60, 0.42)	86312659.55	87054195.44	86262737.85	86179651.78	85604408.06
(0.21, 0.60, 0.62)	87363083.10	87862718.98	87335671.18	87156524.93	86662400.07
(0.21, 0.80, 0.21)	85581492.86	86458359.74	85514990.42	85456543.19	84832720.22
(0.21, 0.80, 0.42)	87037763.86	87773065.20	87002135.10	86894795.69	86314097.68
(0.21, 0.80, 0.62)	88455274.51	88936819.81	88442123.76	88237800.90	87736477.23
(0.35, 0.40, 0.21)	85989751.53	87456069.33	85803578.89	85764525.80	84762967.31
(0.35, 0.40, 0.42)	87196697.64	88448847.67	87054000.56	86935322.61	86004808.81
(0.35, 0.40, 0.62)	88360659.82	89222902.01	88256643.03	87988228.64	87185086.80
(0.35, 0.60, 0.21)	86587737.47	88058158.52	86421159.67	86373511.74	85364265.54
(0.35, 0.60, 0.42)	88419776.00	89671898.02	88308471.20	88182638.99	87237534.36
(0.35, 0.60, 0.62)	90205836.62	91052589.02	90143429.62	89860622.02	89038479.19
(0.35, 0.80, 0.21)	87186255.29	88659682.56	87037216.46	86966200.70	85949535.15
(0.35, 0.80, 0.42)	89654927.59	90902426.04	89571918.35	89421720.31	88462166.14
(0.35, 0.80, 0.62)	92084738.38	92905461.38	92051059.92	91740778.36	90899569.13
(0.48, 0.40, 0.21)	87008789.38	89031364.19	86724031.24	86674342.34	85340786.08
(0.48, 0.40, 0.42)	88685244.24	90420038.82	88461205.62	88307395.54	87058741.51
(0.48, 0.40, 0.62)	90323892.22	91523155.85	90139625.13	89787301.77	88698988.94
(0.48, 0.60, 0.21)	87837704.50	89868788.03	87580827.75	87522755.47	86175341.93
(0.48, 0.60, 0.42)	90390143.80	92130487.16	90212730.44	90051743.23	88776672.34
(0.48, 0.60, 0.62)	92907713.99	94094197.86	92788843.32	92417881.51	91294959.00
(0.48, 0.80, 0.21)	88669323.90	90707516.61	88437453.09	88356707.75	86995769.77
(0.48, 0.80, 0.42)	92120404.25	93859929.57	91985231.13	91799571.19	90498097.12
(0.48, 0.80, 0.62)	95558012.82	96716876.70	95487111.95	95085333.78	93927326.79

References

- [1] W. Avery, G. G. Brown, J. A. Rosenkranz, and R. K. Wood. Optimization of purchase, storage and transmission for natural gas utilities. *Operations Research*, 40:446–462, 1992.
- [2] P. Berling and V. Martínez-de-Albéniz. Optimal inventory policies when purchase price and demand are stochastic. Working Paper, IESE Business School, University of Navarra, Barcelona, Spain, 2009.
- [3] P. Berling and K. Rosling. The effects of financial risk on inventory policy. *Management Science*, 51:1804–1815, 2005.
- [4] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, USA, third edition, 2005.
- [5] J. R. Birge. Option methods for incorporating risk into linear capacity planning models. *Manufacturing & Service Operations Management*, 2:19–31, 2000.
- [6] F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179, 1976.
- [7] A. E. Bopp, V. R. Kannan, S. W. Palocsay, and S. Stevens. An optimization model for planning natural gas purchases, transportation, storage and deliverability. *Omega*, 24:511–522, 1996.
- [8] O. Boyabatli, P. R. Kleindorfer, and S. R. Koontz. Integrating long-term and short-term contracting in fed-cattle supply chains. Working Paper, Lee Kong School of Business, Singapore Management University, Singapore, 2008.
- [9] J. C. Butler and J. S. Dyer. Optimizing natural gas flows with linear programming and scenarios. *Decision Sciences*, 10:563–580, 1999.
- [10] G. M. Constantinides, J. C. Jackwerth, and S. Perrakis. Option pricing: Real and risk-neutral distributions. In J. Birge and V. Linetsky, editors, *Financial Engineering*, Handbooks in Operations Research and Management Science, pages 565–591. Elsevier, Amsterdam, 2007.
- [11] J. C. Cox, J. E. Ingersoll, and S. A. Ross. The relation between forward prices and futures prices. *Journal of Financial Economics*, 9:321–346, 1981.
- [12] J. C. Cox, J. E. Ingersoll, and S. A. Ross. An intertemporal general equilibrium model of asset prices. *Econometrica*, 53:363–384, 1985.

- [13] R. S. Dischel, editor. *Climate Risk and the Weather Market: Financial Risk Management with Weather Hedges*. Risk Books, London, UK, 2002.
- [14] L. Dong and L. Hong. Equilibrium forward contracts on nonstorable commodities in the presence of market power. *Operations Research*, 55:128–145, 2007.
- [15] D. Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ, USA, 2001.
- [16] G. D. Eppen and A. V. Iyer. Improved fashion buying with Bayesian updates. *Operations Research*, 45:805–819, 1997.
- [17] V. Gaur and S. Seshadri. Hedging inventory risk through market instruments. *Manufacturing & Service Operations Management*, 7:103–120, 2005.
- [18] V. Gaur, S. Seshadri, and M. G. Subrahmanyam. Optimal timing of inventory decisions with price uncertainty. New York University, Leonard N. Stern School of Business, New York, NY, USA, 2007.
- [19] S. Gavirneni. Periodic review inventory control with fluctuating purchasing costs. *Operations Research Letters*, 32:374–379, 2004.
- [20] H. Geman. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, Chichester, UK, 2005.
- [21] A. Goel and G. Gutierrez. Integrating commodity markets in the optimal procurement policies of a stochastic inventory system. Working Paper, University of Texas, Austin, TX, USA, 2009.
- [22] S. Graves, H. Meal, S. Dasu, and Y. Qin. Two-stage production planning in a dynamic environment. In S. Axsäter, C. Schneeweiss, and E. Silver, editors, *Multi-Stage Production Planning and Inventory Control*, Lecture Notes in Economics and Mathematical Systems, pages 9–43. Springer-Verlag, Berlin, Germany, 1986.
- [23] J.-M. Guldmann. Supply, storage, and service reliability decisions by gas distribution utilities: A chance-constrained approach. *Management Science*, 29:884–906, 1983.
- [24] J.-M. Guldmann and F. Wang. Optimizing the natural gas supply mix of local distribution utilities. *European Journal of Operations Research*, 112:598–612, 1999.

- [25] H. Gurnani and C. S. Tang. Optimal ordering decisions with uncertain cost and demand forecast updating. *Management Science*, 45:1456–1462, 1999.
- [26] E. G. Haugh. *The Complete Guide to Option Pricing Formulas*. McGraw-Hill, New York, NY, USA, 1998.
- [27] W. H. Hausman. Sequential decision problems: A model to exploit existing forecasters. *Management Science*, 16:B93–B111, 1969.
- [28] W. H. Hausman and R. Peterson. Multiproduct production scheduling for style goods with limited capacity, forecasts revisions, and terminal delivery. *Management Science*, 18:370–383, 1972.
- [29] D. Heath and P. L. Jackson. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions*, 26:17–30, 1994.
- [30] D. P. Heyman and M. J. Sobel. *Stochastic Models in Operations Research: Volume II, Stochastic Optimization*. Dover Publications, Inc., Mineola, NY, USA, 2004.
- [31] A. Huchzermeier and M. A. Cohen. Valuing operational flexibility under exchange rate risk. *Operations Research*, 44:100–113, 1996.
- [32] J. C. Hull. *Options, Futures, and Other Derivatives Securities*. Prentice Hall, Englewood Cliffs, NJ, USA, sixth edition, 2006.
- [33] P. Jaillet, E. I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management Science*, 50:909–921, 2004.
- [34] E. Jouini and H. Kallal. Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory*, 66:178–197, 1995.
- [35] B. A. Kalyon. Stochastic prices in a single-item inventory purchasing model. *Operations Research*, 19:1434–1458, 1971.
- [36] B. Kamrad and S. Lele. Production, operating risk and market uncertainty: A valuation perspective on controlled policies. *IIE Transactions*, 30:455–468, 1998.
- [37] B. Kamrad and P. Ritchken. Valuing fixed price supply contracts. *European Journal of Operational Research*, 74:50–61, 1994.

- [38] P. R. Kleindorfer. Integrating physical and financial risk management in supply management. In H. Geman, editor, *Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy*, pages 33–50. John Wiley & Sons, Chichester, UK, 2008.
- [39] P. R. Kleindorfer and J. Wu. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science*, 49:1597–1615, 2003.
- [40] T. W. Knowles and J. P. Wirick. The Peoples Gas Light and Coke Company plans gas supply. *Interfaces*, 28(5):1–12, 1998.
- [41] B. Kogut and N. Kulatilaka. Operating flexibility, global manufacturing, and the option value of a multinational network. *Management Science*, 40:123–139, 1994.
- [42] P. Kouvelis. Global sourcing strategies under exchange rate uncertainty. In M. M. S. Tayur, R. Ganeshan, editor, *Quantitative Models for Supply Chain Management*, pages 625–668. Kluwer Academic Publishers, Boston, MA, USA, 1998.
- [43] G. Lai, F. Margot, and N. Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research*, Forthcoming, 2009.
- [44] P. J. Lederer and V. R. Singhal. Effect of cost structure and demand risk in the justification of new technologies. *Journal of Manufacturing and Operations Management*, 1:339–371, 1988.
- [45] R. R. Levary and B. V. Dean. A natural gas flow model under uncertainty in demand. *Operations Research*, 28:1360–1374, 1980.
- [46] C.-L. Li and P. Kouvelis. Flexible and risk-sharing supply contracts under price uncertainty. *Management Science*, 45:1378–1398, 1999.
- [47] P. A. Morris, M. J. Sandling, R. B. Fancher, M. A. Kohn, H.-P. Chao, and S. W. Chapel. A utility fuel inventory model. *Operations Research*, 35:169–184, 1987.
- [48] K. Muthuraman, T. Aouam, and R. Rardin. Regulation of natural gas distribution using policy benchmarks. *Operations Research*, 56:1131–1145, 2008.
- [49] J. M. Nascimento and W. B. Powell. An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Mathematics of Operations Research*, 34:210–237, 2009.

- [50] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, New York, NY, USA, 1994.
- [51] P. H. Ritchken and C. S. Tapiero. Contingent claims contracting for purchasing decisions in inventory management. *Operations Research*, 34:864–870, 1986.
- [52] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, USA, 1970.
- [53] E. S. Schwartz. Review of *Investment under Uncertainty* by A. K. Dixit and R. S. Pindyck. *Journal of Finance*, 49:1924–1928, 1994.
- [54] N. Secomandi. Optimal commodity trading with a capacitated storage asset. *Management Science*, Forthcoming, 2009.
- [55] R. W. Seifert, U. W. Thonemann, and W. H. Hausman. Optimal procurement strategies for online spot markets. *European Journal of Operational Research*, 152:781–799, 2004.
- [56] D. Seppi. Risk-neutral stochastic processes for commodity derivative pricing: An introduction and survey. In E. Ronn, editor, *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*, pages 3–60. Risk Publications, London, UK, 2002.
- [57] G. Sick. Real options. In R. A. Jarrow, V. Maksimovic, and W. T. Ziemba, editors, *Finance*, volume 9 of *Handbooks in Operations Research and Management Science*, pages 631–691. Elsevier, Amsterdam, The Netherlands, 1995.
- [58] V. R. Singhal. Inventories, risk, and the value of the firm. *Journal of Manufacturing and Operations Management*, 1:4–43, 1988.
- [59] J. E. Smith. Alternative approaches for solving real-options problems. *Decisions Analysis*, 2:89–102, 2005.
- [60] J. E. Smith and K. F. McCardle. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research*, 47:1–15, 1999.
- [61] N. L. Stokey and R. E. Lucas. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA, USA, 1989. With E. C. Prescott.
- [62] D. M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ, USA, 1998.

- [63] A. J. Triantis and J. E. Hodder. Valuing flexibility as a complex option. *Journal of Finance*, 45:549–565, 1990.
- [64] L. Trigeorgis. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. The MIT Press, Cambridge, MA, USA, 1996.
- [65] J. Williams. *The Economic Function of Futures Markets*. Cambridge University Press, Cambridge, UK, 1986.
- [66] J. Williams. Futures markets: A consequence of risk aversion or transaction costs? *Journal of Political Economy*, 95:1000–1023, 1987.
- [67] J. Yang and Y. Xia. Acquisition management under fluctuating raw material prices. *Production and Operations Management*, 18:212–225, 2009.