

2-2012

Optimal Energy Procurement in Spot and Forward Markets

Nicola Secomandi

Carnegie Mellon University, ns7@andrew.cmu.edu

Sunder Kekre

Carnegie Mellon University, skekre@cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/tepper>



Part of the [Business Administration, Management, and Operations Commons](#)

Optimal Energy Procurement in Spot and Forward Markets

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

Submitted: January 2011; Revised: February 2012

Abstract

Storage capacity for energy, such as electricity, natural gas, and oil, is limited. Thus, spot and forward purchases for delivery on the usage date play an important role in matching the supply and the uncertain demand of energy. Transaction costs tend to be larger in spot than forward energy, and more generally commodity, markets. Hence, partially procuring supply in the forward market, rather than entirely in the spot market, is a potentially valuable real option. We call this option the *forward procurement option*. The study of this option from the perspective of differential transaction costs has received little attention in the literature. We thus formulate and analyze a parsimonious procurement model with differential spot and forward transaction costs and correlated spot demand and nominal price random variables. Our analysis, in part based on natural gas data, sheds novel light on the value of the forward procurement option and its optimal exercise, as well as their sensitivities to parameters of interest. Our main insight is that procuring the demand forecast in the forward market is nearly optimal on the instances that we consider. This greatly simplifies the management of this option. We obtain analogous results with a richer model in which the supply procured in the forward market is delivered at multiple dates. Beyond energy, our research has potential relevance for the procurement of other commodities, such as metals and agricultural products.

1. Introduction

Energy, such as electricity, natural gas, and oil, plays a key economic role in supply chains, as it constitutes a primary input to most industrial and commercial activities (Geman 2005). Spot and forward markets trade energy for immediate and future delivery, respectively. Storage capacity for energy is limited. For example, in 2009 the U.S. natural gas usable storage capacity was about 20% of annual demand (based on data from the Energy Information Administration; EIA 2011). Thus, procurement of supply in spot and forward markets for delivery on the usage date is important in matching a firm's uncertain demand for energy. Indeed, this research was motivated by our collaboration with an energy reselling company that operates without access to storage capacity.

Transaction costs in energy, and more generally commodity, markets tend to be higher for spot than forward trades. Specifically, these costs appear to be inversely related to trading volume, a feature for which there is both theoretical support and empirical evidence (Thompson and Waller 1988, Thompson et al. 1993, Bryant and Haigh 2004, Linnainmaa and Roşu 2009, Roşu 2009). As trading volume typically is thinner far away from the maturity of a futures (forward) contract and during its delivery period, that is, in the spot market, than for dates closer




Figure 1: The average trading volume for the first thirty six maturities of the NYMEX natural gas futures contract during the week of March 8, 2010.

to this maturity, average transaction costs tend to be roughly “U” shaped in time to maturity (see, e.g., Figure 1 in Bryant and Haigh 2004). To illustrate, Figure 1 plots the average trading volume for the first thirty six maturities of the NYMEX natural gas futures contract during the week of March 8, 2010. This figure suggests that trading volume in natural gas futures is higher for contracts whose maturities are closer to the trading date. The reduced trading volume in the spot market reflects the reduced availability of the commodity for discretionary purposes in this market. Consistent with these observations, in their study of oil derivatives, Trolle and Schwartz (2009, p. 4434) state that “open interest for futures contracts tends to peak when expiration is a couple of weeks away, after which open interest declines sharply.”

This feature of transaction costs in energy markets gives firms engaged in the procurement of energy the option to partially satisfy their uncertain commodity requirement (demand) in the forward market, rather than entirely in the spot market. We call this real option the *forward procurement option*. The literature on this topic from the perspective of differential transaction costs is scant (see §2). In this paper, we formulate and analyze a parsimonious model of energy procurement, as well as a richer version of this model, to study the optimal valuation and management of this real option.

In our base model, a firm needs to satisfy an uncertain demand for energy at a given future date (this can be interpreted as a single date or an aggregate of several dates). Energy can

be procured once in the forward market for delivery and usage at this date. Any supply and realized demand mismatch is made up by trading one more final time in the spot market. That is, the firm fully satisfies its demand. Transaction costs are exogenous and proportional to the quantity traded. The spot transaction costs are higher than the forward transaction costs. That is, we model the observed behavior of transaction costs *on average*. In particular, the forward procurement date corresponds to the time when the average forward transaction costs are at their lowest level. The spot energy demand and *nominal* price are correlated random variables (we refer to a price in the absence of transaction costs as a nominal price). We use the valuation approach of Jouini and Kallal (1995), which extends to the case of transaction costs in security trading the risk neutral valuation approach (Luenberger 1998, Chapter 16, Seppi 2002; see Birge 2000 and Smith 2005 for discussions of risk neutral valuation in operations management and decision analysis contexts, respectively). Thus, both in our base and extended models, trading in the forward market is advantageous because it commands lower transaction costs than trading in the spot market.

We show that the size of the optimal forward trade is the optimal solution of a generalization of the classical newsvendor model (Porteus 2002, §1.2) that includes both demand and price, rather than only demand, uncertainty. We characterize the optimality condition of our base model as a generalized critical fractile condition, whereby the optimal weighted overage probability is no more than a critical ratio, which depends on the relevant transaction costs and is always less than $1/2$. This aspect of the critical ratio emphasizes that the key concern in our base model is managing the cost of purchasing too much in the forward market. The value of the forward procurement option is then the difference between the value of the optimal policy of our base model and the value of the policy that only procures in the spot market. This difference can be interpreted as the value of the *cost savings* accrued by optimally trading both in the spot and the forward markets, rather than only in the spot market.

To sharpen our results, we further assume that the spot demand and nominal price are jointly lognormally distributed and satisfy a martingale condition: their expected values, with expectation taken with respect to an equivalent measure in the valuation framework of Jouini and Kallal (1995), are the demand forecast and the forward nominal price, respectively, at the forward trading time. We label this model the martingale lognormal (MLN) model. This model is consistent with the martingale model of demand forecast evolution (Hausman 1969, Graves et al. 1986, Heath and Jackson 1994) and various futures price evolution models, including the one of Black (1976) and the one that ensues when modeling the spot nominal price as a mean

reverting process (see, e.g., Jaillet et al. 2004).

Under the MLN model, we derive closed form expressions for the optimal quantity to procure in the forward market and the value of the forward procurement option. In particular, the optimal forward procurement quantity is the demand forecast scaled by the product of three terms, which we label the covariance, demand variance, and demand riskiness terms. Moreover, we show that the value of the forward procurement option is at most half the spot bid-ask spread. We establish the comparative statics of the value of this option and the optimal forward procurement quantity with respect to several parameters of interest. We also obtain upper bounds on the benefit of optimization, that is, the difference between the values of the optimal forward procurement policy and a policy that procures the demand forecast in the forward market. We use these bounds in the numerical study with our base model.

We apply our closed form expressions in a numerical study in the context of natural gas distribution, based on natural gas data and expert opinions for the transaction costs. We find that the value of the forward procurement option ranges from \$503,364 to \$2,321,028 per month, which amounts to reducing the cost of procuring only in the spot market by 0.75%-3.44%. Despite the magnitude of our computed bounds on the benefit of optimization, we observe that this benefit is small, that is, the demand forecast procurement policy is nearly optimal. This occurs because either the product of the covariance, demand variance, and demand riskiness terms is close to 1 or the objective function of our model is relatively flat around the optimal forward procurement quantity. Consistent with this observation, we observe that both the forward procurement option value and optimal quantity are more sensitive to changes in the standard deviation of the natural logarithm of the spot demand than changes in the natural logarithm of the spot nominal price, as well as changes in the correlation between these random variables.

We also extend our base model in its aggregate delivery date interpretation by disaggregating this date into multiple delivery dates. For example, we disaggregate a monthly time period into days. Our extended model is thus more realistic than our base model. Our structural and numerical analyses of our extended model yield results analogous to the ones obtained with our base model.

Our analysis provides insights into the management of an important business process in energy supply chains. Our insights should be relevant to energy resellers, local distribution companies, and industrial and commercial users, such as food processors, metal and chemical manufacturers, and large restaurant and hotel chains, which purchase large amounts of energy.

Our research suggests that managing the forward procurement option near optimally only requires demand forecasting, without the need to (i) estimate the transaction costs, (ii) model the joint distribution of the spot demand and nominal price, and (iii) optimize the forward procurement decision. Potentially, our insights also have broader applicability for the procurement of other commodities, such as metals and agricultural products.

We proceed by reviewing the related literature in §2. We present our base model in §3. We conduct structural and numerical analyses of this model in §4 and §5, respectively. We formulate and analyze our extended model in §6. We conclude in §7. All the relevant proofs are in Online Appendix A. Online Appendices B-C contain supporting material.

2. Literature Review

Our work is related to the operations management literature on long- and short-term contracting in business-to-business settings (Kleindorfer and Wu 2003, Kleindorfer 2008) and the real option literature on energy and commodity applications (Dixit and Pindyck 1994, Sick 1995, Trigeorgis 1996, Smith and McCardle 1999, Seppi 2002). We add to these literatures by defining the forward procurement option and investigating its optimal exercise and valuation.

Real option models of energy and commodity procurement contracts, such as the swing option and other contracts with varying amounts of sourcing flexibility (see, e.g., Li and Kouvelis 1999, Jaillet et al. 2004), address the uncertainty in the purchase price evolution, but neglect demand uncertainty. In contrast, our models capture the joint uncertainty in the spot demand and nominal price. Our research thus provides novel insights into the effect of demand uncertainty on the optimal forward procurement quantity and the forward procurement option.

The operations management literature includes various models of the uncertain evolution of the demand forecast in procurement (see, e.g., Hausman 1969, Hausman and Peterson 1972, Graves et al. 1986, Heath and Jackson 1994). By integrating a model of demand uncertainty consistent with these demand forecast evolution models and a model of the uncertain evolution of the spot purchase price in a newsvendor setting (Porteus 2002, §1.2), we develop a novel variant of the newsvendor model, for which our analysis generalizes the classical critical fractile condition for the optimal procurement quantity.

Various authors have considered the procurement of a commodity in forward and/or spot markets, including Kalyon (1971), Ritchken and Tapiero (1986), Williams (1986, p. 146), Williams (1987), Gurnani and Tang (1999), Gavirneni (2004), Seifert et al. (2004), Berling and

Rosling (2005), Gaur and Seshadri (2005), Gaur et al. (2007), Goel and Gutierrez (2009), Nascimento and Powell (2009), Oum and Oren (2010), Berling and Martínez-de-Albéniz (2011), and Boyabatli et al. (2011). Different from these authors, we analyze the impact of the correlation between the spot demand and nominal price on the optimal forward procurement quantity and the value of the forward procurement option, and we bring to light the near optimality of the demand forecast forward procurement policy.

Ritchken and Tapiero (1986), Seifert et al. (2004), and Oum and Oren (2010) maximize the expected utility of a risk averse procurement manager, while Gaur and Seshadri (2005), Gaur et al. (2007), and Goel and Gutierrez (2009) model the risk aversion of economic agents via the risk neutral valuation approach (Luenberger 1998, Chapter 16, Seppi 2002). The risk neutral valuation approach does not admit the presence of transaction costs in security trading. In particular, the model of Goel and Gutierrez (2009) includes differential spot and forward transportation costs, while the transaction costs in our models originate from bid-ask spreads, rather than transportation costs. Thus, while we also consider risk averse economic agents, this difference requires us to adopt a valuation framework that admits transaction costs in security trading, specifically the one of Jouini and Kallal (1995). Applying the valuation framework of Jouini and Kallal (1995) involves some discussion of the choice of an equivalent probability measure.

3. Base Model

There is a finite horizon that starts at time 0 and ends at time T . The firm satisfies an energy requirement at time T . We refer to this requirement as the spot demand, denoted by d . The spot demand is a random variable as of any time before time T , and becomes known at this time. In our base model, the time T may correspond to a single date, e.g., a day, or an aggregation of several dates, e.g., a month. Thus, the spot demand is a requirement at a single date, e.g., a daily requirement, in the former case, and an aggregate requirement, e.g., a monthly requirement, in the latter case. Our extended model discussed in §6 refines the aggregate case by disaggregating the aggregate requirement.

The firm does not hold any inventory, due to the lack of access to storage facilities, and can satisfy this requirement by procuring energy in the spot market at time T . The firm also can procure supply in advance at time 0 by entering into a forward contract with physical delivery at time T . In practice, some energy forward contracts can be negotiated to entail delivery on a single date, e.g., a specific day of the month (Belak 2011). When time T is a single

date, our assumption of delivery at time T of the supply procured forward is realistic for such forward contracts. However, most energy forward contracts entail delivery at a constant rate during a given period, e.g., a month. When time T is interpreted as an aggregation of multiple dates, in our base model we do not distinguish when the supply procured forward is actually delivered during the delivery period. In contrast, our extended model discussed in §6 captures the constant and rateable delivery of typical energy forward contracts.

We denote by $F(0)$ the time 0 *nominal* price of a forward contract with time T delivery, and simplify it to F . This price evolves during the time interval $[0, T]$ as a known stochastic process $F(t) \in \mathfrak{R}_+$, $t \in [0, T]$. However, we focus on the time T forward nominal price, which is the spot nominal price f ; that is, $F(T) \equiv f$.

Trading in spot and forward markets incurs transaction costs, that is, bid-ask spreads. Consistent with models studied in the finance literature (see, e.g., Constantinides et al. 2007), we model these costs as proportional. Let A and $B \in (0, 1)$. If the firm purchases spot one unit of energy at time T , it pays the spot ask price $(1 + A)f$; if the firm sells spot one unit of energy at this time, it receives the spot bid price $(1 - A)f$. At time 0 the firm can forward purchase one unit of energy at the forward ask price $(1 + B)F$. We do not allow the firm to short sell energy forward, as this is suboptimal in the no arbitrage valuation framework that we use, which is explained in §4.1. However, a forward sale at time 0 of one unit of energy can be made at price $(1 - B)F$ (this is useful for the discussion in §4.1). Consistent with the theoretical work and the empirical evidence on the structure of transaction costs in energy and commodity markets (see §1), the spot transaction costs are larger than the forward transaction costs; that is, $A > B$.

The firm's time 0 forecast for its time T demand is $D(0)$, which we simplify to D . As the forward price, this forecast may evolve as a stochastic process, denoted by $D(t) \in \mathfrak{R}_+$, $t \in [0, T]$, which is correlated with the forward nominal price stochastic process. However, we focus on the spot demand d at time T ; that is, $D(T) \equiv d$. Since, the demand forecast and forward nominal price stochastic processes are correlated, so are the spot demand and nominal price.

For notational convenience, we define $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | D(t), F(t)]$, $\forall (t, D(t), F(t)) \in [0, T] \times \mathfrak{R}_+^2$, expectation given $D(t)$ and $F(t)$ with respect to a probability distribution discussed in §4.1. Notice that $\mathbb{E}_t[\cdot]$ is not a random variable at time t because $D(t)$ and $F(t)$ are known at this time. That is, our notation does not distinguish between random variables and their realizations; which is which should be clear from the context. We simplify \mathbb{E}_0 to \mathbb{E} .

The firm needs to decide how much supply q to procure forward at time 0. Such an optimal

procurement decision can be obtained by solving the following optimization problem:

$$V := \max_{q \geq 0} \mathbb{E}[(1 - A)f(q - d)^+ - (1 + A)f(q - d)^-] - (1 + B)Fq, \quad (1)$$

where $(\cdot)^+ := \max\{\cdot, 0\}$ and $(\cdot)^- := -\min\{\cdot, 0\}$. The first and second terms inside the expectation in (1) are the revenue collected from selling excess contracted supply on the spot market and the cost of any supply shortfall, respectively. The third term in (1) is the forward procurement cost. The objective function in (1) and its optimal value V are both expressed in time T money, as money is exchanged only at this time.

Due to the presence of the spot market, the firm always achieves 100% service level. Indeed, $q = 0$ is a feasible solution to (1). This solution corresponds to simply waiting until time T , observing the realized demand, and procuring this amount on the spot market. We denote the value of this spot procurement policy as $V^S := -\mathbb{E}[(1 + A)fd]$. We define the value of the forward procurement option, V^P , as the additional value obtained by optimally procuring both in the forward and spot markets than only in the spot market:

$$V^P := V - V^S. \quad (2)$$

4. Structural Analysis of the Base Model

In this section we analyze the optimal exercise and valuation of the forward procurement option by imposing more structure on our base model; in particular, the demand forecast and forward nominal price processes $D(\cdot)$ and $F(\cdot)$, and, hence, the joint probability distribution of the spot demand random variable d and the spot nominal price random variable f . We discuss in §4.1 the valuation framework that we use. We provide a general result in §4.2 and specific results under our MLN model in §4.3.

4.1 Valuation Framework

When trading futures contracts for a commodity commands transaction costs, as we assume, standard risk neutral valuation (Luenberger 1998, Chapter 16, Seppi 2002) of such a commodity cash flows does not apply. We thus apply the valuation framework with transaction costs of Jouini and Kallal (1995) to value these cash flows. These authors show that the absence of arbitrage in securities markets in the presence of transaction costs is equivalent to the existence of at least one probability measure that converts some process between the bid and the ask price processes of traded securities into a martingale. Any such measure, which Jouini and

Kallal (1995) refer to as an equivalent measure or, alternatively, a martingale measure, can be used for valuation purposes in a manner analogous to risk neutral valuation.

To apply this theory here, suppose that the forward transaction costs do not decrease in the time interval $[0, T]$. That is, let $B(\cdot)$ be a time dependent function such that $B(0) \equiv B$, $B(t) \leq B(t')$, $\forall t, t' \in [0, T]$, $t \leq t'$, and $B(T) \equiv A$. Refer to the process $F(t)$, $t \in [0, T]$, as the process for the *nominal* price of a futures contract for time T delivery; that is, when the transaction costs are absent (here we assume equivalence between futures and forward nominal prices; this condition holds when the risk free interest rate is deterministic, Cox et al. 1981, which we also assume to be the case here). Then, a probability measure under which the process $F(\cdot)$ is a martingale is an equivalent measure in the valuation framework of Jouini and Kallal (1995). That is, under this measure it holds that $F(t) = \mathbb{E}_t[F(t')]$, $\forall t, t' \in [0, T]$, $t \leq t'$, and $F(t) \in \mathfrak{R}_+$.

We use this martingale measure to determine the equivalent, according to Jouini and Kallal (1995), probability distribution of the spot nominal price random variable. This is a natural choice of equivalent measure. For example, in the absence of transaction costs it is the only choice consistent with risk neutral valuation. That is, this equivalent measure is a risk adjusted measure for the futures nominal price process. Under our chosen equivalent measure, the forward nominal price F is thus equal to the expected spot nominal price f (Shreve 2004, p. 244):

$$F = \mathbb{E}[f]. \tag{3}$$

We assume that changes in the demand forecast process $D(\cdot)$ are uncorrelated with changes in the price of the market portfolio. Hence, we do not risk adjust the evolution of this stochastic process (Smith 2005). (This does not imply that the demand forecast and the forward nominal price processes are uncorrelated.) This is realistic for commercial and residential energy demand, as it is largely determined by the weather (see the related discussion in Hull 2012, Chapter 33).

We focus on the uncertainty in the spot demand and nominal price given the information available at the beginning of the time horizon, that is, D and F . All probabilistic statements related to these random variables are under their joint distribution as determined by the previous assumptions. For simplicity, we assume that spot demand is a continuous random variable.

Condition (3) implies that the expected spot nominal price conditional on the time 0 forward price is finite (as $F \in \mathfrak{R}_+$). Assumption 1 further states that the conditional expectation of the product of the spot demand and nominal price is finite.

Assumption 1 (Bounded spot nominal price and demand product expectation). $\mathbb{E}[fd] < \infty$, $\forall(D, F) \in \mathfrak{R}_+^2$.

Condition (3) and Assumption 1 imply that $\mathbb{E}[d] < \infty$, $\forall(D, F) \in \mathfrak{R}_+^2$, because they hold when the random variables f and d are independent, in which case we have that $\infty > \mathbb{E}[fd] = \mathbb{E}[f]\mathbb{E}[d] = F\mathbb{E}[d]$, so that $\mathbb{E}[d] < \infty$.

4.2 A General Result

Let $\mathbb{C}(f, d)$ denote the covariance between f and d conditional on D and F . The value of the spot procurement policy, V^S , can be easily expressed as

$$V^S = -(1 + A)\{\mathbb{E}[f]\mathbb{E}[d] + \mathbb{C}(f, d)\} = -(1 + A)\{F\mathbb{E}[d] + \mathbb{C}(f, d)\}. \quad (4)$$

Determining the value of the optimal forward procurement policy, $V(D, F)$, is more involved. Proposition 1 characterizes the optimality condition for problem (1) without assuming a specific (equivalent) joint probability distribution for the spot demand and nominal price.

Proposition 1 (Optimality condition). *Assume that $F > 0$. Consider a positive feasible solution to problem (1). Such a solution is optimal if and only if satisfies the following condition:*

$$\mathbb{E} \left[\frac{f}{F} \mathbb{P}\{d \leq q \mid f\} \right] = \frac{1}{2} \left(1 - \frac{B}{A} \right). \quad (5)$$

If no positive solution to problem (1) satisfies this condition, then zero is the optimal solution to this problem.

As shown in the proof of Proposition 1, the objective function of problem (1) can be expressed as

$$(A - B)Fq - 2A\mathbb{E}[f(q - d)^+] - \mathbb{E}[(1 + A)fd]. \quad (6)$$

The third term in (6) is the value of procuring only in the spot market. The first term in (6) is the value saved by the firm by purchasing an amount of supply q at time 0, irrespective of whether this amount is needed to satisfy the demand at time T . The second term in (6) is the reduction in this value if q exceeds this demand. Expression (6) thus shows that problem (1) can be thought of as maximizing the value of the net savings from operating both in the forward and spot markets relative to only operating in the spot market. In other words, the central concern in problem (1) is balancing the savings from procuring in the forward market, rather than only in the spot market, against the cost of buying too much in advance.

To obtain more insight into this trade off, it is useful to derive condition (5) using marginal analysis. If a positive quantity is procured at time 0, then its associated marginal underage and overage costs are $(1 + A)f - (1 + B)F$ and $(1 + B)F - (1 - A)f$, respectively. Intuitively, such a decision is optimal, in which case we denote it by q^* , if and only if it balances its expected marginal underage and overage costs. Let $1\{\cdot\}$ be the indicator function of event $\{\cdot\}$. Exploiting condition (3), the optimal expected marginal underage cost is

$$\begin{aligned}\mathbb{E}[(1 + A)f - (1 + B)F]1\{d > q^*\} &= \mathbb{E}[(1 + A)f - (1 + B)F](1 - 1\{d \leq q^*\}) \\ &= (A - B)F - \mathbb{E}[(1 + A)f - (1 + B)F]1\{d \leq q^*\}.\end{aligned}$$

The optimal expected marginal overage cost is $\mathbb{E}[(1 + B)F - (1 - A)f]1\{d \leq q^*\}$. We obtain (5) by equating the expressions for these costs, rearranging, and noticing that $\mathbb{E}[f1\{d \leq q^*\}] = \mathbb{E}[f\mathbb{E}[1\{d \leq q^*|f\}]] = \mathbb{E}[f\mathbb{P}\{d \leq q^*|f\}]$.

We can further interpret condition (5) as a generalization of the classical newsvendor critical fractile optimality condition (Porteus 2002, §1.2), which relates the overage probability and the critical ratio. Consider the right hand side of (5). The expectations of the marginal underage and overage costs are $(A - B)F$ and $(A + B)F$, respectively. Thus, the term $(1 - B/A)/2 \in (0, 1/2)$ is the ratio of the expected marginal underage cost and the sum of the expected marginal underage and overage costs, and can be interpreted as a critical ratio. The term f/F , with $F > 0$, is nonnegative and such that $\mathbb{E}[f/F] = 1$, so that it acts as a weighting random variable for the term $\mathbb{P}\{d \leq q^* | f\}$. Hence, the left hand side of (5) assumes the interpretation of weighted overage probability. Condition (5) relates the weighted overage probability and the critical ratio at optimality. Given this interpretation, that the right hand side of (5) is positive and less than 1/2 emphasizes that the key concern in problem (1) is managing the cost component associated with purchasing too much supply in advance.

4.3 Results with the MLN Model

In general, condition (5) can be used as the basis for developing numerical methods to compute an optimal forward procurement decision at time 0 and, hence, the value of the optimal procurement policy. By making specific assumptions on the joint probability distribution of the spot demand and nominal price, we can obtain closed form expressions for these quantities.

We proceed by assuming a lognormal model with unbiased demand forecast. Specifically, we assume that the random variables spot demand d and spot nominal price f are jointly lognormally distributed, and $D = \mathbb{E}[d]$. That is, we assume that the natural logarithms of the

Table 1: The time variance, covariance, and riskiness terms.

Term	Notation	Definition
Covariance	K^C	$\exp(cs_d s_f)$
Demand Variance	K^V	$\exp(-s_d^2/2)$
Demand Riskiness	K^R	$\exp(zs_d)$

random variables d and f are jointly normally distributed with standard deviations s_d and s_f and means $\ln D - s_d^2/2$ and $\ln F - s_f^2/2$, respectively, and correlation coefficient c . This is our MLN model. It satisfies condition (3) and Assumption 1.

The MLN model is consistent with a version of the single product multiplicative martingale model of demand forecast evolution of Heath and Jackson (1994), which in our case essentially reduces to the model of Hausman (1969). It is also consistent with the futures price evolution model of Black (1976), as well as the one implied by assuming that the spot nominal price follows a mean reverting process (see, e.g., Jaillet et al. 2004).

The optimal forward procurement quantity. Proposition 2 establishes a closed form expression for the optimal forward procurement quantity under the MLN model. This proposition uses the covariance, demand variance, and demand riskiness terms defined in Table 1. The covariance term depends on the covariance between the random variables $\ln d$ and $\ln f$; that is, $\rho s_d s_f$. The demand variance term depends on the variance of the random variable $\ln d$; that is, s_d^2 . We denote by z the $(1 - B/A)/2$ -th percentile of the standard normal distribution. The demand riskiness term depends on the product of z , which is negative as the critical ratio $(1 - B/A)/2$ is positive and less than $1/2$, and the standard deviation of $\ln d$ (the riskiness label is motivated by the negative sign of the exponent of this term and contrasts the safety label that would be appropriate if this sign were positive).

Proposition 2 (Optimal forward procurement quantity with the MLN model). *Under the MLN model the optimal amount of supply to procure forward in problem (1) is*

$$q^* = K^C K^V K^R D. \quad (7)$$

This proposition shows that under the MLN model the optimal quantity to procure forward is the current demand forecast D scaled by the product of the covariance, demand variance, and demand riskiness terms. In isolation, the demand variance and riskiness terms decrease the optimal amount of supply to procure in the forward market below the demand forecast. The effect of the covariance term depends on the sign of the correlation coefficient c . In isolation, this effect is to make q^* decrease below D if $c < 0$ and increase above D if $c > 0$. It is easy to

Table 2: Comparative statics of the optimal forward procurement quantity under the MLN model; $-$, 0 , and $+$ denote a negative, null, and positive effect, respectively.

Parameter	Sign of the Effect
D	$+$
c	$+$
s_f	$c < 0$: $-$; $c = 0$: 0 ; $c > 0$: $+$
s_d	$c \leq 0$: $-$; $c > 0$: $+$ for $s_d \in (0, \bar{s}_d^q)$, $-$ for $s_d \in (\bar{s}_d^q, \infty)$
B/A	$-$

see that q^* is at least D if and only if $c \geq (s_d/2 - z)/s_f$. A sufficiently large value of c is thus needed for the optimal procurement quantity to exceed D .

Comparative statics of the optimal forward procurement quantity. Under the MLN model the optimal forward procurement quantity does not depend on the forward price. It depends only on the standard deviation of the natural logarithm (\log) of the spot nominal price and the correlation coefficient between the \log spot demand and nominal price. Corollary 1, which follows easily from Proposition 2, establishes the comparative statics of the optimal forward procurement quantity with respect to other quantities of interest. We define $\bar{s}_d^q := cs_f + z$. Table 2 summarizes these comparative statics.

Corollary 1 (Comparative statics of q^* with the MLN model). *Consider the MLN model. Ceteris paribus, the optimal forward procurement quantity (1) increases in D ; (2) increases in c ; (3) decreases in s_f if $c < 0$, does not depend on s_f if $c = 0$, and increases in s_f if $c > 0$; (4) decreases in s_d if $c \leq 0$, and increases in s_d when s_d increases up to \bar{s}_d^q and decreases in s_d for values of this parameter that exceed \bar{s}_d^q if $c > 0$; and (5) decreases in B/A .*

The increase of the optimal forward procurement quantity in the demand forecast, D , is obvious as this amounts to an increase in the expected spot demand.

The behavior of this optimal quantity in the correlation coefficient, c , is less obvious, but can be explained as follows. Under the MLN model, the probability that the spot demand and the spot nominal price exceed or are less than or equal to the given pair (\bar{d}, \bar{f}) increases in c (this holds because of Theorem 11 in Müller 2001, and (9.A.18) and the discussion preceding (9.A.17) in Shaked and Shanthikumar 2007, p. 395). As the spot demand and the spot nominal price are both bounded below from zero but are not bounded above by any pair of real numbers, procuring more in the forward market becomes more appealing as c increases.

The effect of the log spot nominal price standard deviation, s_f , depends on the sign of the correlation coefficient c and modulates its effect. An increase in s_f weakens the effect of c if

$c < 0$, it strengthens it if $c > 0$, and does not affect it if $c = 0$.

Changing the log spot demand standard deviation, s_d , affects the demand variance, demand riskiness, and covariance terms. If $c \leq 0$, an increase in s_d affects all of these terms in the same direction, so that the optimal forward procurement quantity decreases. If $c > 0$, the increase in the covariance term due to an increase in s_d dominates the corresponding decreases in the demand variance and riskiness terms, and the optimal procurement quantity also increases; but when s_d grows sufficiently large, the decreases in the latter two terms dominate the increase in the covariance term, and the optimal procurement quantity also decreases.

The decrease of the optimal forward procurement quantity in the transaction cost ratio, B/A , is intuitive.

Valuations. Proposition 3 provides closed form expressions for the values of the spot procurement policy, the optimal forward procurement policy, and the forward procurement option, as well as a bound on the relative value of this option. Two of these expressions are closed form in the sense that they depend on the cumulative distribution function of the standard normal distribution, $\Phi(\cdot)$, which can be readily evaluated in a spreadsheet, through the term $\alpha := 2A\Phi(z - s_d)$.

Proposition 3 (Valuations with the MLN model). *Consider model MLN. The values of the spot procurement policy, the optimal procurement policy, the forward procurement option, and a bound on the relative value of this option, respectively, are*

$$V^S = -(1 + A)K^C FD, \quad (8)$$

$$V = V^S + \alpha K^C FD, \quad (9)$$

$$V^P = \alpha K^C FD, \quad (10)$$

$$\frac{V^P}{|V^S|} < A. \quad (11)$$

The value of the spot procurement policy is the product of the transaction cost term $-(1 + A)$, the covariance term, the forward nominal price, and the demand forecast. The value of the optimal procurement policy is the value of the spot procurement policy plus the value of the net savings that accrue to the firm by optimally procuring in the forward market, that is, the term $\alpha K^C FD$ (see the discussion following expression (6)). By definition (2), the value of these net savings is the value of the forward procurement option. Obviously, the value of this option cannot be negative. Moreover, this value relative to the absolute value of the spot procurement option is bounded above by half the spot market bid-ask spread, A . That is, this is the most

Table 3: Comparative statics of the value of the forward procurement option with the MLN model; $-$, 0 , and $+$ denote a negative, null, and positive effect, respectively.

Parameter	Sign of the Effect
D	$+$
F	$+$
c	$+$
s_f	$c < 0$: $-$; $c = 0$: 0 ; $c > 0$: $+$
s_d	either $c \leq 0$ or $c > 0$ and C holds: $-$ $c > 0$ and \bar{C} holds: $+$ for $s_d \in (0, \bar{s}_d^V)$, $-$ for $s_d \in (\bar{s}_d^V, \infty)$
B/A	$-$

that can be gained from optimally procuring both in the forward and the spot markets, rather than only in the spot market.

Comparative statics of the value of the forward procurement option. Corollary 2 establishes some of the comparative statics of the value of the forward procurement option. We denote by $\phi(\cdot)$ the standard normal density function. We abbreviate by C the condition $s_f c \leq \phi(z)/\Phi(z)$, and by \bar{C} its complement. If \bar{C} is true, then we define \bar{s}_d^V as that value of s_d that satisfies the equality $s_f c = \phi(z - s_d)/\Phi(z - s_d)$ (\bar{s}_d^V exists and is positive as shown in the proof of Corollary 2).

Corollary 2 (Comparative statics of V^P with the MLN model). *Consider model MLN. Ceteris paribus, the value of the forward procurement option, V^P , (1) increases in D ; (2) increases in F ; (3) increases in c ; (4) decreases in s_f if $c < 0$, is not affected by s_f if $c = 0$, and increases in s_f if $c > 0$; (5) decreases in s_d if either $c \leq 0$ or $c > 0$ and C hold, and increases up to the threshold \bar{s}_d^V and decreases thereafter in s_d if $c > 0$ and \bar{C} holds; and (6) decreases in B/A .*

Consider part (1). From expression (6), the value of the forward procurement option is the difference between the value of the savings and the value of the overprocurement cost when acting optimally in the forward market. These values are both linear in the forward nominal price, F (this is obvious for the savings; see (39) in Appendix A for the overprocurement cost). The value of these savings is obviously greater than the value of this overprocurement cost for every given such price. The difference between the former and the latter values thus increases in this price. A similar argument holds for part (2) with respect to the spot demand forecast, D . Parts (3)-(4) mirror how the correlation coefficient and the log spot nominal price standard deviation affect the optimal forward procurement action. Part (5) is related to the behavior of the optimal forward procurement quantity in the log spot demand standard deviation. Part (6) is intuitive.

The benefit of optimization. Determining the optimal forward procurement quantity requires knowledge of the parameters of the joint spot demand and nominal price distribution, as well as the spot and forward market transaction costs. It is thus of interest to assess the benefit of optimization by investigating the suboptimality of a forward procurement policy that does not require this information.

We take this policy to be the one that procures forward an amount equal to the demand forecast, that is, $q = D$. We choose this policy as it does not require knowledge of the transaction costs and ignores all but one of the parameters of the MLN model; that is, D . We denote by V^F the value of this policy (expression (42) in Online Appendix B provides a formula for computing this quantity). The absolute benefit of optimization is the difference $V - V^F$. The relative benefit of optimization is the ratio $(V - V^F)/|V|$.

The benefit of optimization is zero when there is no *demand* uncertainty (irrespective of the specific spot nominal price distribution used). This is because the choice in this case is between procuring an amount D in the forward market if $(1 + B)F \leq (1 + A)\mathbb{E}[f] = (1 + A)F$ and procuring this amount in the spot market otherwise, and $B < A$. We thus expect optimization to yield small benefits for small values of the log spot demand standard deviation parameter s_d , which under the MLN model is related to the coefficient of variation (CV) of spot demand via $\sqrt{\exp(s_d^2) - 1}$. However, it is not clear what one should expect this benefit to be for larger values of this parameter.

We can obtain more insights into the benefit of optimization by studying bounds on its performance. As shown in Online Appendix C, the absolute and relative benefits of optimization under the MLN model can be bounded as follows:

$$V - V^F < \begin{cases} (2A + B)FD, & \text{if } c < 0, \\ 3AFD, & \text{if } c = 0, \\ 3AK^C FD, & \text{if } c > 0, \end{cases} \quad (12)$$

$$\frac{V - V^F}{|V|} < \begin{cases} (2A + B)/K^C, & \text{if } c < 0, \\ 3A, & \text{if } c \geq 0. \end{cases} \quad (13)$$

The bounds in (12) suggest that optimization is potentially more beneficial in absolute terms when the correlation between the log spot demand and nominal price is nonnegative (since $A > B$ and $K^C > 1$ if $c > 0$; as discussed in §5.1, the case of positive correlation is the more realistic one). However, the bounds in (13) suggest that the opposite may be true in relative terms for large values of the log spot demand and/or nominal price standard deviations (since K^C should be small in this case, as $K^C \in (0, 1)$ if $c < 0$). Indeed, it can be verified that

in this case V , V^F , and $(V - V^F)/|V|$, respectively, approach 0, $-(A + B)FD$, and ∞ if $c < 0$; $-(1 + A)FD$, $-(1 + 2A + B)FD$, and $(A + B)/(1 + A)$ if $c = 0$; and $-\infty$, $-\infty$, and 0 if $c > 0$.

In particular, for the more realistic case of positively correlated log spot demand and nominal price, even though the absolute benefit of optimization can grow large when the parameters s_d and/or s_f do so, in relative terms this benefit is bounded. In §5 we numerically investigate the benefit of optimization in a realistic setting, in which the log spot demand and nominal prices are positively correlated.

5. Numerical Study with the Base Model

In this section we report the results of a numerical study with our base model aimed at assessing the value of the forward procurement option and how to capture this value, as well as illustrating some of the comparative statics of this option and its optimal exercise. We present the instances used as test beds in §5.1. We discuss our results in §5.2.

5.1 Instances

In our study, the firm is a hypothetical natural gas distributor operating in the U.S. We let the length of the time horizon be two weeks, and thus set T equal to 14/365. The end of the time horizon corresponds to February 2011; that is, we use the aggregate interpretation of our base model. We take the firm's February 2011 demand forecast as of time 0 to be 14,593,766MMBtu, which corresponds to the February 2001 demand faced by the utility considered in the study of Muthuraman et al. (2008, Table 1, p. 1143). Here, we make the convenient assumption that this figure is indicative of the demand forecast of this firm for February 2011. We use \$4.4315/MMBtu as the forward nominal price at the beginning of the time horizon; this is the average of the low and high prices of the NYMEX natural gas futures contract for delivery in February 2011 observed on January 14, 2011.

We specify the parameters of the MLN model by employing the single product multiplicative martingale model of demand forecast evolution of Heath and Jackson (1994) and the seasonal mean reverting model for energy spot prices discussed by Jaillet et al. (2004). Specifically, the demand forecast for time T 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. The log deseasonalized natural gas spot nominal price is $\chi(t) := \ln(f(t)/S(t))$, where $f(t)$ is the spot nominal price at time t and $S(t)$ is the deterministic seasonality factor for this price. This log price evolves as a mean reverting

process with speed of mean reversion κ , *risk adjusted* mean reversion level ξ , and volatility σ_χ . The dynamics of these processes in the time interval $[0, T]$ are

$$dD(t) = \sigma_D D(t) dZ_D(t), \quad (14)$$

$$d\chi(t) = \kappa[\xi - \chi(t)]dt + \sigma_\chi dZ_\chi(t), \quad (15)$$

$$dZ_D(t)dZ_\chi(t) = \rho dt, \quad (16)$$

where $dZ_D(t)$ and $dZ_\chi(t)$ are increments to standard Brownian motions with instantaneous correlation coefficient ρ , and dt is an infinitesimal time increment. Under model (14)-(16), the parameters of the MLN model are as follows:

$$\begin{aligned} s_d &= \sigma_D \sqrt{T}, \\ s_f &= \sqrt{\frac{\sigma_\chi^2}{2\kappa}(1 - e^{-2\kappa T})}, \\ c &= \frac{\rho[(1 - e^{-\kappa T})/\kappa]}{\sqrt{T} \sqrt{(1 - e^{-2\kappa T})/(2\kappa)}}. \end{aligned}$$

Given that F is known, the values of the parameters $\chi(0)$, ξ , and $S(T)$ are not needed, because they are assumed to satisfy condition (3), which reduces to

$$F = S(T) \exp \left(\chi(0)e^{-\kappa T} + (1 - e^{-\kappa T})\xi + \frac{\sigma_\chi^2}{4\kappa}(1 - e^{-2\kappa T}) \right).$$

However, the values of these parameters are needed when dealing with our extended model, as discussed in §6.3.

We set a base value for the demand forecast volatility, σ_D , to make the CV of spot demand equal to 0.05, which seems a reasonable figure for natural gas demand. In our specification of the MLN model, the CV of spot demand given the information available at time 0 is $\sqrt{\exp(\sigma_D^2 T) - 1}$. Thus, for given values of T and this CV, the corresponding value of σ_D is $\sqrt{\ln(1 + \text{CV}^2)/T}$. For T equal to 14/365 and the CV of spot demand equal to 0.05, we obtain a value for σ_D equal to 0.2551, which we round to 0.26. To obtain insights into the effect of this parameter on the value of the forward procurement option, we also consider values for σ_D equal to 0.51, 0.76, 1.01, 1.26, and 1.50, which roughly correspond to spot demand CV values equal to 0.10, 0.15, 0.20, 0.25, and 0.30, respectively.

For the log deseasonalized spot nominal price speed of mean reversion and volatility, κ and σ_χ , we consider base values of 1.0547 and 0.6696, which are the estimates of these parameters obtained by Lai et al. (2011, Table 2) using NYMEX data. Similar to σ_D , we consider additional

Table 4: Parameter values used in our numerical study with our base model; when there are multiple values for a parameter, the boldface value corresponds to the base case.

Parameter	Value(s)
T	14/365
D	14,593,766MMBtu
F	\$4.475/MMBtu
σ_D	0.26 , 0.51, 0.76, 1.01, 1.26, 1.50
σ_χ	0.2696, 0.3696, 0.4696, 0.5696, 0.6696 , 0.7696
ρ	0.1, 0.2 , 0.3, 0.4, 0.5, 0.6
A	3.75%
B	0.025% , 0.25%, 2.5%

values for σ_χ , namely 0.2696, 0.3696, 0.4696, 0.5696, and 0.7696, leaving the value of the parameter κ at its base level.

We set a base value for the instantaneous correlation coefficient, ρ , by making the correlation between the spot demand and nominal price with $T = 14/365$ equal to 0.2 (this value is used in the study of Seifert et al. 2004). Specifically, in our specification of the MLN model the correlation between the spot demand and nominal price is $\{\exp(\rho\sigma_D\sigma_\chi[1 - \exp(-\kappa T)]/\kappa) - 1\}/(G_D G_f)$, with $G_D := \sqrt{\exp(\sigma_D^2 T) - 1}$ and $G_f := \sqrt{\exp(\sigma_\chi^2 [1 - \exp(-2\kappa T)]/(2\kappa)) - 1}$. Thus, for given values of this correlation, denoted by CORR, and T , the corresponding value of ρ is $\ln(1 + G_D G_f \text{CORR})/\{\sigma_D \sigma_\chi [1 - \exp(-\kappa T)]/\kappa\}$. For $T = 14/365$, $\sigma_D = 0.26$, and $\sigma_\chi = 0.6696$, we obtain $\rho = 0.1968$, which we round to 0.2. Based on the argument made by Seifert et al. (2004) that the correlation between a commodity price and the demand for this commodity faced by a firm should be positive, we also consider values for ρ equal to 0.1, 0.3, 0.4, 0.5, and 0.6, each of which roughly corresponds to an equal CORR value.

In terms of the transaction costs, we rely on indicative figures provided to us by practitioners at a major U.S. energy trading company and at a major U.S. natural gas broker. These figures are consistent with our assumed transaction cost structure. We set the forward bid-ask spread ($2B$) equal to 0.05%, and the spot bid-ask spread ($2A$) to 7.5%. We also consider forward bid-ask spreads equal to 0.5% and 5% to quantify the impact of the forward and spot transaction cost differential on the forward procurement option value and optimal quantity.

Table 4 summarizes the parameters used to generate our 648 ($= 6^3 \cdot 3$) instances.

5.2 Results

In discussing our results, we focus on the value of the forward procurement option, the optimal forward procurement quantity, some of their comparative statics, and the benefit of optimiza-

Table 5: Summary of our numerical results with the base model.

Quantity	Base Case	Range
V^P	\$2,338,924	\$503,364-2,321,028
$V^P/ V^S $	3.44%	0.75-3.44%
q^*/D	99.96%	72.19-100.30%
$V - V^F$	\$3	\$0-258,130
$(V - V^F)/ V $	0.00%	0.00-0.39%

tion. Table 5 summarizes our results.

The value of the forward procurement option. The upper bound (11) on the relative value of the forward procurement option is 3.75%. In the base case, the relative value of this option is 3.44%, which is 91.73% of this bound. Thus, the bound (11) is informative in this case. Across all our instances, the relative value of the forward procurement option varies in between 0.75% and 3.44%, which correspond to 20.00% and 91.73% of the value of the bound (11). In particular, the ratios of the relative value of the forward procurement option and this bound vary in between 73.60% and 91.73%, 68.00% and 86.13%, and 20.00% and 29.87% when the forward transaction cost is 0.025%, 0.25%, and 2.5%, respectively. The bound (11) thus becomes looser when the forward transaction cost increases.

In absolute terms, the forward procurement option is worth \$2,338,924 per month in the base case, and ranges from \$503,364 to \$2,321,028 per month across all our instances. Since these are reductions in the cost of the spot procurement policy achieved by the optimal forward procurement policy, these findings suggest that a firm may derive substantial benefit from optimally exploiting the differential transaction costs between the forward and spot markets.

Comparative statics of the value of the forward procurement option. Proposition 3 implies that the ratio of the forward procurement option and the absolute value of the spot procurement policy is equal to $\alpha/(1+A)$, which only depends on the transaction cost parameters, A and B , the demand forecast volatility, σ_D , and the length of the time horizon, T (as α depends on the product $\sigma_D T$ via $s_d = \sigma_D \sqrt{T}$, and z , which in turn depends on A and B). Increasing the demand forecast volatility decreases these ratios by less than 1%. Specifically, these ratios vary in between 2.76% and 3.44%, 2.55% and 3.23%, and 0.75% and 1.12% when the forward transaction cost is 0.025%, 0.25%, and 2.5%, respectively. As expected, each set of ratios decreases when the forward transaction cost increases. More interesting, each set of ratios decreases below 2% only when this cost increases by two orders of magnitude, that is, from 0.025% to 2.5%. Further, even when the spot and forward transaction costs share the same order of magnitude, the forward procurement option retains some value.




Figure 2: The value of the forward procurement option as a function of the demand forecast volatility, σ_D , for different values of the instantaneous correlation coefficient, ρ ; the log deseasonalized spot nominal price volatility and the forward transaction cost are at their base levels, $\sigma_\chi = 0.6696$ and $B = 0.025\%$; the base case values of σ_D and ρ are 0.35 and 0.1, respectively.

Even though these ratios do not vary by more than 1%, for a given value of the forward transaction cost, the monetary value of the forward procurement option varies in a more substantial manner. Specifically, it ranges in between \$1,852,171 and \$2,321,028, \$1,714,462 and \$2,175,388, and \$503,364 and \$751,599 per month when the forward transaction cost is 0.025%, 0.25%, and 2.5%, respectively. These variations in the option value are almost completely attributable to changes in the demand forecast volatility, σ_D (increasing σ_D decreases the option value in our instances). That is, the value of the forward procurement option is rather insensitive to changes in the log deseasonalized spot nominal price volatility, σ_χ , and the instantaneous correlation coefficient, ρ , in our instances.

For illustration purposes, Figure 2 shows the value of the forward procurement option as a function of the demand forecast volatility, σ_D , for different values of the instantaneous correlation coefficient, ρ , with the log deseasonalized spot nominal price volatility and the forward procurement cost at their base levels, $\sigma_\chi = 0.60$ and $B = 0.025\%$. Figure 3 displays the value of the forward procurement option as a function of the demand forecast volatility, σ_D , for different values of the log deseasonalized spot nominal price volatility, σ_χ , with the




Figure 3: The value of the forward procurement option as a function of the demand forecast volatility, σ_D , for different values of the log deseasonalized spot nominal price volatility, σ_χ ; the instantaneous correlation coefficient and the forward transaction cost are at their base levels, $\rho = 0.1$ and $B = 0.025\%$; the base case values of σ_D and σ_χ are 0.26 and 0.6696, respectively.

instantaneous correlation coefficient and the forward procurement cost at their base levels, $\rho = 0.10$ and $B = 0.025\%$. The predominant effect of the demand forecast volatility on the value of the forward procurement option relative to the two other parameters is evident in these figures.

The optimal forward procurement quantity and its comparative statics. The optimal forward procurement quantity, q^* , is 99.96% of the demand forecast, D , in the base case. It varies in between 72.19% and 100.30% of this forecast across all our instances. More specifically, it ranges in between 95.69% and 100.30%, 93.59% and 99.90%, and 72.19% and 95.50% of the demand forecast when the forward transaction cost is 0.025%, 0.25%, and 2.5%, respectively. The optimal forward procurement quantity thus drops below 93% of the demand forecast in some of the instances in which the spot and forward transaction costs are of the same order of magnitude.

To explain this observation, recall from Proposition 2, specifically (7), that the optimal forward procurement quantity is the demand forecast scaled by the product of the demand variance, demand riskiness, and covariance terms. In isolation, the demand variance and riski-

ness terms make the optimal forward procurement quantity decrease below the demand forecast, while the covariance term has the opposite effect when the instantaneous correlation coefficient is positive, which is the case in our numerical study. Across all our instances, the covariance and demand variance terms range from 1.0003 to 1.0264 and from 0.9578 to 0.9987, respectively, and their product ranges from 0.9592 to 1.0039. Thus, the combined effect of these two terms is to keep the optimal forward procurement quantity close to the demand forecast. In contrast to these two terms, the demand riskiness term depends on the forward transaction cost. This term varies in between 0.9975 and 0.9996, 0.9757 and 0.9957, and 0.7526 and 0.9519, when this cost is 0.025%, 0.25%, and 2.5%, respectively. Consequently, the optimal forward procurement quantity drops considerably below the demand forecast only in some instances when the spot and forward transaction costs share the same order of magnitude.

Similar to the value of the forward procurement option, the optimal forward procurement quantity is more sensitive to the volatility of the demand forecast than to the volatility of the log deseasonalized spot nominal price and the instantaneous correlation coefficient.

The benefit of optimization. The values of the upper bound (12) on the absolute benefit of optimization range from \$194,067,951 to \$199,134,740 per month across all of our instances; the value of the upper bound (13) on the relative benefit of optimization is 11.25%. However, when the spot and forward transaction costs have different orders of magnitude, the analysis of the optimal forward procurement quantity suggests that there may be little benefit from optimizing the forward procurement quantity relative to setting it equal to the demand forecast. Indeed, both the absolute and the relative benefits of optimization are small, being no more than \$258,130 per month and 0.39%, respectively, across all our instances. Thus, the bounds (12) and (13) are loose on our instances.

Given our discussion of the optimal forward procurement quantity, the finding that there is little benefit from optimization is expected when the spot and forward transaction costs have different orders of magnitude, as in this case the demand forecast is close to the optimal forward quantity, but otherwise it is more surprising. In the latter case, even though the demand forecast can be substantially above the optimal forward procurement quantity, that is, by as much as 32.87% ($= (1/0.7526 - 1)\%$) of this quantity, we have observed that the objective function of problem (1) is relatively flat around its optimal solution. This finding suggests that achieving near optimal management of the forward procurement option only requires demand forecasting, rather than estimating both the transaction costs, the joint distribution of the spot demand and nominal price, and optimizing the forward procurement quantity. Thus, these results suggest

that the simple policy that procures in the forward market the demand forecast is appealing for practical implementation.

6. Extended Model and Its Analysis

This section extends our base model in its aggregation of dates interpretation by disaggregating these dates (§6.1); partially characterizes the optimal policy of this extended model and provides a formula to compute its value under an extension of the MLN model (§6.2); and conducts a numerical study of this policy and its value similar to the one performed with our base model (§6.3).

6.1 Formulation

In our extended model, the aggregate time period T is subdivided into I equally spaced dates, each denoted by time T_i with $i \in \mathcal{I} := \{1, \dots, I\}$. For example, if the time period T represents a month, then I is the number of days in this month and the date T_i is the i -th day in this month. There is still a single purchase at time 0 in the forward market. Purchasing an amount of supply q in the forward market at this time entails a payment equal to $(1 + B)Fq$ at time T_1 (recall that F is the forward nominal price at time 0). Different from the base model, a fraction $1/I$ of the quantity q purchased at time 0 is delivered on each date T_i . That is, our base model corresponds to the special case in which $I = 1$.

We continue to use an equivalent measure under which the forward price process $F(\cdot)$ is a martingale during the time interval $(0, T_1]$. Since there are now multiple cash flows on different dates during the delivery period, we explicitly model the risk free discount factor. We denote by δ the risk free discount factor from time T_i back to time T_{i-1} for each $i \in \mathcal{I} \setminus \{1\}$. We let f_i be the spot nominal price at time T_i . We extend the spot and forward nominal price convergence condition stated in §3, that is, $F(T) \equiv f$. Specifically, at the beginning of the delivery period, time T_1 , the forward nominal price converges to the fraction $1/I$ of the sum of the discounted expected values of the spot nominal prices during the delivery period:

$$F(T_1) \equiv \frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}_1[f_i], \quad (17)$$

where \mathbb{E}_1 is expectation conditional on the information available at time T_1 .

Demand is date specific. The spot demand on date T_i is the quantity d_i . At time 0, the firm's forecast of the *total* demand on dates T_1 through T_I is D . The firm's time 0 forecast of

the demand on date T_i is the fraction β_i of the total demand forecast; that is,

$$\mathbb{E}[d_i | D] = \beta_i D, \quad (18)$$

with $\beta_i \geq 0$ for all $i \in \mathcal{I}$ and $\sum_{i \in \mathcal{I}} \beta_i = 1$. The firm must fully satisfy demand on each date T_i . Due to the lack of storage, the firm can do this either by employing the amount of the quantity procured forward and delivered on each date, q/I , or by procuring in the spot market on the date. Excess supply delivered on a given date is sold into the spot market. That is, our extended model includes multiple spot make up transactions. For simplicity, we assume that the spot transaction costs are constant across the dates T_1 through T_I . That is, the purchase and sale spot prices on date T_i are $(1 + A)f_i$ and $(1 - A)f_i$, respectively. Each spot demand and nominal price are correlated random variables. Analogous to our base model, we assume that each spot demand random variable is a continuous and nonnegative random variable.

The firm must decide at time 0 the quantity to purchase forward for delivery at times T_1 through T_I . That is, at time 0 the firm's needs to solve the following optimization problem:

$$V := \max_{q \geq 0} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E} \left[(1 - A)f_i \left(\frac{q}{I} - d_i \right)^+ - (1 + A)f_i \left(\frac{q}{I} - d_i \right)^- \right] - (1 + B)Fq. \quad (19)$$

This optimization problem assumes that all cash flows are valued as of time T_1 . Problem (19) reduces to problem (1) when $I = 1$.

The spot procurement policy in this extended setting corresponds to procuring the realized spot demand on each date in the delivery period. The value of this policy is thus $V^S = -\sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}[(1 + A)f_i d_i]$. The value of the forward procurement option V^P remains the difference $V - V^S$.

6.2 Structural Analysis

We focus on characterizing the structure of the optimal forward procurement policy and its value. We start this analysis by pointing out that condition (17) implies that the forward nominal price at time 0, F , satisfies

$$F = \frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}[f_i]. \quad (20)$$

This is true because

$$F = \mathbb{E}[F(T)] = \mathbb{E} \left[\frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}_1[f_i] \right] = \frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}[\mathbb{E}_1[f_i]] = \frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}[f_i],$$

where the first equality is due to our use of an equivalent measure, under which the forward nominal price is a martingale, the second equality follows from condition (17), and the last equality holds by the law of iterated expectations.

We use condition (20) in Proposition 4 to characterize the optimality condition of problem (19). To make this characterization more specific, we also extend our MLN model to the setting of this section. We assume that the log spot demand and nominal price on each date T_i are jointly normally distributed with respective standard deviations $s_{d,i}$ and $s_{f,i}$, respective means $\ln(\beta_i D) - s_{d,i}^2/2$ and m_i , and correlation coefficient c_i . This model satisfies condition (18). We also assume that the parameters m_i 's and F satisfy condition (20).

Although in general it is impossible to obtain a closed form expression for the optimal solution to problem (19) under our extended MLN model, the value of such an optimal solution under this model can be expressed in closed form. To this aim, we extend the terms K^C and α to the terms $K_i^C := \exp(c_i s_{d,i} s_{f,i})$ and $\alpha_i(q) := 2A\Phi(\ln((q/I)/(\beta_i D))/s_{d,i} - c_i s_{f,i} - s_{d,i}/2)$. Proposition 4 provides partial characterizations of the optimal forward procurement quantity, and closed form expressions for the value of the spot procurement policy, the optimal forward procurement policy, and the forward procurement option under our extended MLN model.

Proposition 4 (Extended model). (1) *A positive solution to problem (19) is optimal if and only if it satisfies the condition*

$$\sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E} \left[\frac{f_i}{F} \mathbb{P} \left\{ d_i \leq \frac{q}{I} \mid f_i \right\} \right] = \frac{1}{2} \left(1 - \frac{B}{A} \right). \quad (21)$$

If no positive solution satisfies this equation, then zero optimally solve problem (19).

(2) *Under the extended MLN model, the optimal solution to problem (19) is positive and satisfies the condition*

$$\sum_{i \in \mathcal{I}} \left\{ \frac{\delta^{i-1} \mathbb{E}[f_i]}{F} \right\} \Phi \left(\frac{\ln((q/I)/(\beta_i D))}{s_{d,i}} - c_i s_{f,i} + \frac{s_{d,i}}{2} \right) = \frac{1}{2} \left(1 - \frac{B}{A} \right). \quad (22)$$

(3) *Moreover, the values of the spot procurement policy, the optimal procurement policy, and the forward procurement option are*

$$V^S = -(1+A)D \sum_{i \in \mathcal{I}} \delta^{i-1} K_i^C \beta_i \mathbb{E}[f_i], \quad (23)$$

$$V = V^S + \sum_{i \in \mathcal{I}} \alpha_i(q^*) K_i^C \beta_i D \delta^{i-1} \mathbb{E}[f_i], \quad (24)$$

$$V^P = \sum_{i \in \mathcal{I}} \alpha_i(q^*) K_i^C \beta_i D \delta^{i-1} \mathbb{E}[f_i]. \quad (25)$$

Proposition 4 generalizes Proposition 1 and partially Propositions 2 and 3. Extending the first order optimality condition (5), condition (21) includes in its left hand side the sum of the discounted weighted overage probabilities, one for each date in the delivery period, and in its right hand side the product of the critical ratio $(1 - B/A)/2$. Under our extended MLN model, this first order optimality condition reduces to the condition (22). In general, this condition does not admit a closed form expression, as is instead possible with the base model under the MLN model, because each term in the sum in its left hand side involves the standard normal cumulative distribution evaluated at a different values. The expressions (23), (24), and (25) for the values of the spot procurement policy, the optimal forward procurement policy, and the forward procurement option are qualitatively consistent with their respective analogous expressions (8), (9), and (10) under the MLN model.

As in our base model, we continue to measure the benefit of optimization as the difference between the values of the optimal forward procurement policy and the demand forecast forward procurement policy. Expression (44) in Online Appendix B provides a formula for computing the value of the latter policy. We use this formula in §6.3.

6.3 Numerical Study

We extend our numerical study discussed in §5 by taking the time interval $[T_1, T_I]$ to be the month of February 2011. We interpret the parameter I as the number of days in this month, and thus set it equal to 28. We let the risk free discount rate to be 0.01 and define $\delta := \exp(-0.01/365)$.

We specify the extended MLN model in a manner similar to our specification of the MLN model in §5.1. We continue to model the evolution of the log deseasonalized spot nominal price $\chi(t)$ as a mean reverting model. As in Lai et al. (2011, Table 2), we set the parameter ξ equal to -2.0421 and each parameter $S(T_i)$ equal to 1.0761 (the February seasonality parameter, a constant). We then use condition (20) to determine the spot nominal price $f(0)$. That is, under model (15) the mean and standard deviations of the time T_i log deseasonalized spot nominal price $\chi(T_i)$ satisfy, respectively,

$$\begin{aligned} m_i &= \chi_0 e^{-\kappa T_i} + \xi(1 - e^{-\kappa T_i}), \\ s_{f,i} &= \sqrt{\frac{\sigma_\chi^2}{2\kappa}(1 - e^{-2\kappa T_i})}. \end{aligned}$$

It follows that

$$\mathbb{E}[f_i | f(0)] = S(T_i) \exp \left(\chi(0)e^{-\kappa T_i} + \xi(1 - e^{-\kappa T_i}) + \frac{\sigma^2}{4\kappa}(1 - e^{-2\kappa T_i}) \right). \quad (26)$$

For a given set of parameters, we thus solve the equation

$$F = \frac{1}{I} \sum_{i \in \mathcal{I}} \delta^{i-1} S(T_i) \exp \left(\chi(0)e^{-\kappa T_i} + \xi(1 - e^{-\kappa T_i}) + \frac{\sigma^2}{4\kappa}(1 - e^{-2\kappa T_i}) \right) \quad (27)$$

in the unknown $\chi(0)$. Equation (27) always has a unique solution, as its right hand side tends to 0 when $\chi(0)$ tends to $-\infty$ and it monotonically increases in $\chi(0)$.

We obtain the distribution of the spot demand for time T_i by modeling the evolution of the *forecast* of this demand, which we denote by $D_i(\cdot)$, in a manner analogous to how we model the evolution of the aggregate demand forecast $D(\cdot)$ in §5.1. Specifically, we let each $D_i(\cdot)$ evolve during the time interval $[0, T_i]$ as a driftless geometric Brownian motion correlated with the log deseasonalized spot nominal price process. That is, for all $i \in \mathcal{I}$ we assume that

$$dD_i(t) = \sigma_i D_i(t) dZ_i(t), \quad (28)$$

$$dZ_i(t) dZ_\chi(t) = \rho_i dt, \quad (29)$$

where σ_i is the volatility of the spot demand forecast for date T_i , $dZ_i(t)$ is an increment to a standard Brownian motion specific to date T_i , and ρ_i is the instantaneous correlation coefficient between this standard Brownian motion and the one associated with the log deseasonalized spot nominal price in (15).

Under model (15) and (28)-(29), the standard deviation of the log spot demand is

$$s_{d,i} = \sigma_i \sqrt{T_i},$$

and the correlation between this random variable and the log deseasonalized spot nominal price is

$$c_i = \frac{\rho_i [(1 - e^{-\kappa T_i}) / \kappa]}{\sqrt{T_i} \sqrt{(1 - e^{-2\kappa T_i}) / (2\kappa)}}.$$

For simplicity, we set the numerical values of each of the parameters σ_i and ρ_i equal to the ones of the parameters σ_D and ρ considered in §5.1. Also for simplicity, we assume that each parameter β_i is equal to $1/I = 1/28$.

Our numerical results with the extended model, summarized in Table 6, mirror the ones obtained with the base model. Specifically, the value of the forward procurement option is \$2,279,846 per month in the base case and ranges from \$410,709 to \$2,293,523 per month across

Table 6: Summary of our numerical results with the extended model.

Quantity	Base Case	Range
V^P	\$2,279,846	\$410,709-2,293,523
$V^P/ V^S $	3.39%	0.61-3.39%
q^*/D	99.94%	62.83-100.56%
$V - V^F$	\$5	\$0-369,030
$(V - V^F)/ V $	0.00%	0.00-0.55%

all the considered instances. These figures correspond to improvements of 3.39% and 0.61-3.39%, respectively, on the cost of the spot procurement policy. In the base case, the optimal forward procurement quantity is 99.94% of the demand forecast D , that is, the demand forecast for the entire time period $[T_1, T_I]$. Across all our instances, the optimal forward procurement quantity varies in between 62.83% and 100.56% of this forecast. This range is 92.37-100.56%, 89.67-100.00%, and 62.83-94.11% when the forward transaction cost is 0.025%, 0.25%, and 2.5%, respectively. The demand forecast forward procurement policy is near optimal, as the benefit of optimization is no more than 0.55%, that is, \$369,030 per month. These results suggest that the insights into the value of the forward procurement option, the relationship between the optimal and the demand forecast forward procurement policies, and the benefit of optimization obtained with our base model have broader applicability, as they carry through to the setting of our extended model.

7. Conclusions

In this paper we study the optimal valuation and management of the forward procurement option, a real option that arises from the differential transaction costs in spot and forward energy, and more broadly commodity, markets. We do this first by formulating and analyzing a parsimonious procurement model with correlated spot demand and nominal price random variables. We characterize the optimal solution of our model. Based on further specific assumptions on the spot demand and nominal price random variables, we derive closed form expressions for the optimal forward procurement quantity and the value of the forward procurement option, also establishing several of their comparative statics, and bounds on the value of this option and the benefit of optimization. In a numerical study based on natural gas data and expert opinions, we investigate the value of the forward procurement option and its optimal exercise, illustrate some of their comparative statics, and quantify the benefit of optimization. Our main insight from this study is that managing the forward procurement option by procuring the

demand forecast in the forward market is near optimal on the instances that we consider. We also formulate and analyze a richer procurement model in which the supply procured in the forward market is delivered on multiple dates, and obtain analogous results. Our analysis and results provide novel managerial insights into the management of a fundamental business process in energy supply chains. Our findings are of particular significance to energy resellers and local distribution companies, but retain potential relevance in other industrial and commercial contexts.

Our work can be extended in several directions. We consider a single procurement date in the forward market. It would be of interest to consider a more dynamic model with additional procurement dates later on in the forward market. This model may yield some additional benefit relative to our models, but would require estimating the dynamics of the transaction costs in a more granular fashion.

Firms may face limits to the amount of trading that they can perform in the forward market at a given date, e.g., due to limited working capital availability, counterparty risk when transacting in the forward market, as well as supply chain disruptions (Tomlin and Wang 2012) that may affect the availability of contracted supply. Our models could be extended to consider these features.

The transaction costs are deterministic in our models. This simplification captures the average behavior of these costs. A more refined approach might model the stochastic behavior of transaction costs, possibly linking it to the one of trading volume, and the relationship between the evolution of these costs and the ones of demand and price. Relevant empirical research also would be useful in this respect.

We model a single date when demand is to be satisfied. With multiple usage dates, storage (inventory) considerations would become important, despite the limited availability of energy storage capacity in practice. Our models could be extended in this direction, in a manner similar to the inventory models reviewed in §2 but also by modeling transaction costs and the joint evolution of demand and price. In particular, it would be of interest to assess the value of storage for energy users in practical settings. Relevant practice based work in this area includes the stochastic programming approach used by The Peoples Gas Light and Coke Company to plan its natural gas supply (Knowles and Wirick 1998; see also the models and applications discussed by Guldmann 1983, Morris et al. 1987, Bopp et al. 1996, Butler and Dyer 1999, Guldmann and Wang 1999); the deterministic optimization approach described by Avery et al. (1992) and applied at Southwest Gas Corporation and Questar Pipeline Corporation; and

the optimization approach of Levary and Dean (1980) that periodically updates the relevant demand forecast in a reactive fashion and revises its earlier supply decisions accordingly, and has been applied to data pertaining to the East Ohio Gas Company. Integrating our approach, which includes transaction costs and the joint evolution of demand and price, with the practical aspects of these models may offer additional managerial insights.

Acknowledgments

We thank the members of the review team, the constructive feedback of which led to an improved version of this paper. Nicola Secomandi was supported by NSF grant CMMI1129163.

References

- Avery, W., G. G. Brown, J. A. Rosenkranz, R. K. Wood. 1992. Optimization of purchase, storage and transmission for natural gas utilities. *Operations Research* **40**(3) 446–462.
- Belak, R. 2011. Personal communication.
- Berling, P., V. Martínez-de-Albéniz. 2011. Optimal inventory policies when purchase price and demand are stochastic. *Operations Research* **59**(1) 109–124.
- Berling, P., K. Rosling. 2005. The effects of financial risk on inventory policy. *Management Science* **51**(12) 1804–1815.
- Birge, J. R. 2000. Option methods for incorporating risk into linear capacity planning models. *Manufacturing & Service Operations Management* **2**(1) 19–31.
- Black, F. 1976. The pricing of commodity contracts. *Journal of Financial Economics* **3**(1-2) 167–179.
- Bopp, A. E., V. R. Kannan, S. W. Palocsay, S. P. Stevens. 1996. An optimization model for planning natural gas purchases, transportation, storage and deliverability. *Omega* **24**(5) 511–522.
- Boyabatli, O., P. R. Kleindorfer, S. R. Koontz. 2011. Integrating long-term and short-term contracting in beef supply chains. *Management Science* **57**(10) 1771–1787.
- Bryant, H. L., M. S. Haigh. 2004. Bid-ask spreads in commodity futures markets. *Applied Financial Economics* **14**(13) 923–936.
- Butler, J. C., J. S. Dyer. 1999. Optimizing natural gas flows with linear programming and scenarios. *Decision Sciences* **30**(2) 563–580.
- Constantinides, G. M., J. C. Jackwerth, S. Perrakis. 2007. Option pricing: Real and risk-neutral distributions. J. Birge, V. Linetsky, eds., *Financial Engineering*. Handbooks in Operations Research and Management Science, Elsevier, Amsterdam, 565–591.
- Cox, J. C., J. E. Ingersoll, S. A. Ross. 1981. The relation between forward prices and futures prices. *Journal of Financial Economics* **9**(4) 321–346.
- Dixit, A. K., R. S. Pindyck. 1994. *Investment under Uncertainty*. Princeton University Press, Princeton, NJ, USA.
- EIA. 2011. Natural gas. http://www.eia.gov/dnav/ng/ng_stor_cap_dcu_nus_a.htm. Accessed on 1/22/2011.
- Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing & Service Operations Management* **7**(2) 103–120.
- Gaur, V., S. Seshadri, M. G. Subrahmanyam. 2007. Optimal timing of inventory decisions with price uncertainty. New York University, Leonard N. Stern School of Business, New York, NY, USA.

- Gavirneni, S. 2004. Periodic review inventory control with fluctuating purchasing costs. *Operations Research Letters* **32**(4) 374–379.
- Geman, H. 2005. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, Chichester, UK.
- Goel, A., G. Gutierrez. 2009. Integrating commodity markets in the optimal procurement policies of a stochastic inventory system. Working Paper, University of Texas, Austin, TX, USA.
- Graves, S., H. Meal, S. Dasu, Y. Qin. 1986. Two-stage production planning in a dynamic environment. S. Axsäter, C. Schneeweiss, E. Silver, eds., *Multi-Stage Production Planning and Inventory Control*. Lecture Notes in Economics and Mathematical Systems, Springer-Verlag, Berlin, Germany, 9–43.
- Guldmann, J.-M. 1983. Supply, storage, and service reliability decisions by gas distribution utilities: A chance-constrained approach. *Management Science* **29**(8) 884–906.
- Guldmann, J.-M., F. Wang. 1999. Optimizing the natural gas supply mix of local distribution utilities. *European Journal of Operations Research* **112**(3) 598–612.
- Gurnani, H., C. S. Tang. 1999. Optimal ordering decisions with uncertain cost and demand forecast updating. *Management Science* **45**(10) 1456–1462.
- Hausman, W. H. 1969. Sequential decision problems: A model to exploit existing forecasters. *Management Science* **16**(2) B93–B111.
- Hausman, W. H., R. Peterson. 1972. Multiproduct production scheduling for style goods with limited capacity, forecasts revisions, and terminal delivery. *Management Science* **18**(7) 370–383.
- Heath, D., P. L. Jackson. 1994. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions* **26**(3) 17–30.
- Hull, J. C. 2012. *Options, Futures, and Other Derivatives Securities*. Eight ed. Prentice Hall, Englewood Cliffs, NJ, USA.
- Jaillet, P., E. I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management Science* **50**(7) 909–921.
- Jouini, E., H. Kallal. 1995. Martingales and arbitrage in securities markets with transaction costs. *Journal of Economic Theory* **66**(1) 178–197.
- Kalymon, B. A. 1971. Stochastic prices in a single-item inventory purchasing model. *Operations Research* **19**(6) 1434–1458.
- Kleindorfer, P. R. 2008. Integrating physical and financial risk management in supply management. H. Geman, ed., *Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy*. John Wiley & Sons, Chichester, UK, 33–50.
- Kleindorfer, P. R., D. J. Wu. 2003. Integrating long- and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science* **49**(11) 1597–1615.
- Knowles, T. W., J. P. Wirrick. 1998. The Peoples Gas Light and Coke Company plans gas supply. *Interfaces* **28**(5) 1–12.
- Lai, G., M. X. Wang, S. Kekre, A. Scheller-Wolf, N. Secomandi. 2011. Valuation of storage at a liquefied natural gas terminal. *Operations Research* **59**(3) 602–616.
- Levary, R. R., B. V. Dean. 1980. A natural gas flow model under uncertainty in demand. *Operations Research* **28**(6) 1360–1374.
- Li, C.-L., P. Kouvelis. 1999. Flexible and risk-sharing supply contracts under price uncertainty. *Management Science* **45**(10) 1378–1398.
- Linnainmaa, J. T., I. Roşu. 2009. Weather and time series determinants of liquidity in a limit order market. Working Paper, Booth School of Business, University of Chicago, Chicago, IL, USA.
- Luenberger, D. G. 1998. *Investment Science*. Oxford University Press, New York, NY, USA.
- Morris, P. A., M. J. Sandling, R. B. Fancher, M. A. Kohn, H.-P. Chao, S. W. Chapel. 1987. A utility fuel inventory model. *Operations Research* **35**(2) 169–184.
- Müller, A. 2001. Stochastic ordering of multivariate normal distributions. *Annals of the Institute of Statistical Mathematics* **53**(3) 567–575.

- Muthuraman, K., T. Aouam, R. Rardin. 2008. Regulation of natural gas distribution using policy benchmarks. *Operations Research* **56**(6) 1131–1145.
- Nascimento, J. M., W. B. Powell. 2009. An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Mathematics of Operations Research* **34**(1) 210–237.
- Oum, Y., S. S. Oren. 2010. Optimal static hedging of volumetric risk in a competitive wholesale electricity market. *Decision Analysis* **7**(1) 107–122.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford Business Books, Stanford, CA, USA.
- Ritchken, P. H., C. S. Tapiero. 1986. Contingent claims contracting for purchasing decisions in inventory management. *Operations Research* **34**(6) 864–870.
- Roşu, I. 2009. A dynamic model of the limit order book. *Review of Financial Studies* **22**(11) 4601–4641.
- Seifert, R. W., W. T. Ulrich, H. H. Warren. 2004. Optimal procurement strategies for online spot markets. *European Journal of Operational Research* **152**(3) 781–799.
- Seppi, D. 2002. Risk-neutral stochastic processes for commodity derivative pricing: An introduction and survey. E. Ronn, ed., *Real Options and Energy Management Using Options Methodology to Enhance Capital Budgeting Decisions*. Risk Publications, London, UK, 3–60.
- Shaked, M., J. G. Shanthikumar. 2007. *Stochastic Orders*. Springer, NY, NY, USA.
- Shreve, S. E. 2004. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer, New York, NY, USA.
- Sick, G. 1995. Real options. R. A. Jarrow, V. Maksimovic, W. T. Ziemba, eds., *Finance, Handbooks in Operations Research and Management Science*, vol. 9. Elsevier, Amsterdam, The Netherlands, 631–691.
- Smith, J. E. 2005. Alternative approaches for solving real-options problems. *Decisions Analysis* **2**(2) 89–102.
- Smith, J. E., K. F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Operations Research* **47**(1) 1–15.
- Thompson, S., J. Eales, D. Seibold. 1993. Comparison of liquidity costs between Kansas City and Chicago wheat futures contracts. *Journal of Agricultural and Resource Economics* **18**(2) 185–197.
- Thompson, S., M. Waller. 1988. Determinants of liquidity costs in commodity futures markets. *Review of Futures Markets* **7**(1) 110–126.
- Tomlin, B., Y. Wang. 2012. Operational strategies for managing supply chain disruption risk. P. Kouvelis, L. Dong, O. Boyabatli, R. Li, eds., *Handbook of Integrated Risk Management in Global Supply Chains*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 79–101.
- Trigeorgis, L. 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. The MIT Press, Cambridge, MA, USA.
- Trolle, A. B., E. S. Schwartz. 2009. Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies* **22**(11) 4423–4461.
- Williams, J. 1986. *The Economic Function of Futures Markets*. Cambridge University Press, Cambridge, UK.
- Williams, J. 1987. Futures markets: A consequence of risk aversion or transaction costs? *Journal of Political Economy* **95**(5) 1000–1023.

Online Appendix

A. Proofs

Proof of Proposition 1 (Optimality condition). Define $v(q) := (1 - A)f(q - d)^+ - (1 + A)f(q - d)^-$. Since $(\cdot)^- \equiv (-\cdot)^+ = -\cdot + (\cdot)^+$, it holds that $v(q; d, f) = (1 + A)f(q - d) - 2Af(q - d)^+$. This and condition (3) imply that the objective function of problem (1) can be rearranged as

$$\begin{aligned} \mathbb{E}[v(q)] - (1 + B)Fq &= \mathbb{E}[(1 + A)f(q - d) - 2Af(q - d)^+] - (1 + B)Fq \\ &= (A - B)Fq - 2A\mathbb{E}[f(q - d)^+] - \mathbb{E}[(1 + A)fd]. \end{aligned}$$

Thus, solving problem (1) is equivalent to solving

$$\max_{q \geq 0} (A - B)Fq - 2A\mathbb{E}[f(q - d)^+]. \quad (30)$$

Given condition (3) and Assumption 1, it is easy to show that the objective function of problem (30) is proper concave in q (see Rockafellar 1970, p. 24 for the definition of proper concave function), and tends to $-\infty$ as q tends to ∞ . Thus, a positive solution to this problem is optimal if and only if it satisfies the first order optimality condition

$$(A - B)F - 2A\mathbb{E}[f1\{d \leq q\}] = 0, \quad (31)$$

where $1\{\cdot\}$ is the indicator function of event $\{\cdot\}$. The expectation in (31) can be expressed as $\mathbb{E}[f1\{d \leq q\}] = \mathbb{E}[f\mathbb{E}[1\{d \leq q\}|f]] = \mathbb{E}[f\mathbb{P}\{d \leq q|f\}]$. Using this equality in (31), rearranging, and exploiting the assumed positivity of F yields

$$\mathbb{E}\left[\frac{f}{F}\mathbb{P}\{d \leq q|f\}\right] = \frac{1}{2}\left(1 - \frac{B}{A}\right).$$

If no positive solution satisfies this condition, the concavity in q of the objective function of problem (30) implies that zero is an optimal solution to this problem. \square

Lemma 1 (Lognormal partial moment). *Denote by $f(w; \mu, \sigma)$ and $\phi(y; \mu, \sigma)$ the probability density functions of a lognormal random variable W and a normal random variable Y , respectively, with parameters μ and σ , and by $\Phi(\cdot)$ the cumulative distribution function of the standard normal random variable. Let $\gamma \in \mathfrak{R}$ and $q \in \mathfrak{R}_+$. It holds that*

$$\mathbb{E}[W^\gamma 1\{W \leq q\}] = \int_0^q w^\gamma f(w; \mu, \sigma) dw = e^{(\frac{\sigma^2 \gamma}{2} + \mu)\gamma} \Phi((\ln q - (\mu + \sigma^2 \gamma))/\sigma).$$

Proof. It holds that

$$\begin{aligned} \int_0^q w^\gamma f(w; \mu, \sigma) dw &= \int_0^q e^{\gamma \ln w} f(w; \mu, \sigma) dw = \int_{-\infty}^{\ln q} e^{\gamma y} \phi(y; \mu, \sigma) dy \\ &= \int_{-\infty}^{\ln q} \frac{1}{\sigma \sqrt{2\pi}} e^{\gamma y - \frac{(y-\mu)^2}{2\sigma^2}} dy, \end{aligned} \quad (32)$$

where the second equality follows from $f(w; \mu, \sigma) = \phi(\ln w; \mu, \sigma)/w$ and the change of variable $y = \ln w$. By adding and subtracting $\sigma^4\gamma^2 + 2\sigma^2\mu\gamma$ to the exponent of the exponential term in (32), this exponent can be written as $[(\sigma^2\gamma)/2 + \mu]\gamma - [y - (\sigma^2\gamma + \mu)]^2/(2\sigma^2)$. Substituting this expression for this exponent yields

$$\int_{-\infty}^{\ln q} \frac{1}{\sigma \sqrt{2\pi}} e^{\gamma y - \frac{(y-\mu)^2}{2\sigma^2}} dy = e^{(\frac{\sigma^2\gamma}{2} + \mu)\gamma} \int_{-\infty}^{\ln q} \phi(y; \mu + \sigma^2\gamma, \sigma) dy = e^{(\frac{\sigma^2\gamma}{2} + \mu)\gamma} \Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right). \square$$

Proof of Proposition 2 (Optimal forward procurement quantity with the MLN model). The expectation in (31) can be expressed as $\mathbb{E}[f1\{d \leq q\}] = \mathbb{E}[1\{d \leq q\}\mathbb{E}[f|d]]$. The properties of the bivariate lognormal distribution imply that under the MLN model it holds that

$$\mathbb{E}[f|d] = Fd^{cs_f/s_d} D^{-cs_f/s_d} e^{cs_f(s_d - cs_f)/2}.$$

This expression implies that

$$\mathbb{E}[1\{d \leq q\}\mathbb{E}[f|d]] = FD^{-cs_f/s_d} e^{cs_f(s_d - cs_f)/2} \mathbb{E}[d^{cs_f/s_d} 1\{d \leq q\}]. \quad (33)$$

By Lemma 1 applied to the lognormal random variable $d|D$ with parameters $\mu := \ln D - s_d^2/2$, $\sigma := s_d$, and $\gamma := cs_f/s_d$, it follows that

$$\mathbb{E}[d^{cs_f/s_d} 1\{d \leq q\}] = D^{cs_f/s_d} e^{-cs_f(s_d - cs_f)/2} \Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right).$$

Hence, (33) can be written as

$$\mathbb{E}[1\{d \leq q\}\mathbb{E}[f|d]] = F\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right). \quad (34)$$

Since $\mathbb{E}[1\{d \leq q\}\mathbb{E}[f|d]] = \mathbb{E}[s\mathbb{P}\{d \leq q|f\}]$, substituting the right hand side of (34) into the left hand side of (5) yields

$$\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right) = \frac{1}{2} - \frac{B}{2A}, \quad (35)$$

which implies $[\ln q - (\mu + \sigma^2\gamma)]/\sigma = z$, or, equivalently, $q = D \exp(-s_d^2/2 + cs_d s_f + z s_d)$, that is, (7) (once q is replaced with q^*). \square .

Lemma 2 is useful to establish Proposition 3.

Lemma 2 (Expected overage value with the MLN model). *Consider the MLN model. Let $\mu := \ln D - s_d^2/2$, $\sigma := s_d$, and $\gamma := cs_f/s_d$. Given $q > 0$, it holds that*

$$\mathbb{E}[f(q-d)^+] = F \left[q\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right) - De^{cs_dsf}\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma} - \sigma\right) \right].$$

Proof. Analogous to the proof of Proposition 2 (see the derivation of (33)), it holds that

$$\mathbb{E}[f(q-d)^+] = \mathbb{E}[\mathbb{E}[f|d](q-d)^+] = FD^{-\gamma}e^{cs_f(s_d-cs_f)/2}\mathbb{E}[d^\gamma(q-d)1\{d \leq q\}]. \quad (36)$$

Expanding the last expectation, applying Lemma 1 twice, and rearranging terms yields

$$\begin{aligned} \mathbb{E}[d^\gamma(q-d)1\{d \leq q\}] &= q\mathbb{E}[d^\gamma 1\{d \leq q\}] - \mathbb{E}[d^{\gamma+1}1\{d \leq q\}] \\ &= e^{(\sigma^2\gamma/2+\mu)\gamma}q\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right) \\ &\quad - e^{(\sigma^2\gamma/2+\mu)\gamma}e^{\mu+\sigma^2/2+\sigma^2\gamma}\Phi\left(\frac{\ln q - (\mu + \sigma^2(\gamma+1))}{\sigma}\right) \\ &= D^\gamma e^{-cs_f(s_d-cs_f)/2} \\ &\quad \left[q\Phi\left(\frac{\ln q - (\mu + \sigma^2\gamma)}{\sigma}\right) - De^{cs_dsf}\Phi\left(\frac{\ln q - (\mu + \sigma^2(\gamma+1))}{\sigma}\right) \right]. \end{aligned}$$

Substituting the last expression into (36) yields the claimed result. \square

Proof of Proposition 3 (Valuations with the MLN model). The properties of the bivariate lognormal distribution imply that under the MLN model it holds that

$$\mathbb{E}[fd] = \mathbb{E}[\exp(\ln f + \ln d)] = FDK^C. \quad (37)$$

Expression (8) follows from (37) by accounting for the transaction cost of operating in the spot market.

At optimality, the objective function of problem (1) is

$$V = (A-B)Fq^* - 2A\mathbb{E}[f(q^*-d)^+] - (1+A)\mathbb{E}[fd]. \quad (38)$$

Applying Lemma 2 with $q = q^*$ and recalling that $\Phi(z) = 1/2 - B/(2A)$ yields

$$\mathbb{E}[f(q^*-d)^+] = F \left[\left(\frac{1}{2} - \frac{B}{2A}\right)q^* - DK^C\Phi(z - \sigma_D\sqrt{T}) \right]. \quad (39)$$

Expression (9) follows from substituting (39) into (38), and expressing the term $\mathbb{E}[fd]$ by (37).

Together with (2), (9) implies (10).

The inequality (11) holds because

$$\frac{V^P}{|V^S|} = \frac{2A\Phi(z-s_d)K^CDF}{(1+A)K^CDF} = \frac{2A\Phi(z-s_d)}{1+A} < A,$$

where the last step follows from $\Phi(z-s_d) \in (0, 0.5)$ and $A > 0$. \square

Proof of Corollary 2 (Comparative statics of V^P with the MLN model). Parts (1)-(4) and (6) follow easily from (10). Consider part (5). The sign of the partial derivative of V^P with respect to s_d is the sign of

$$s_f c \Phi(z - s_d) - \phi(z - s_d). \quad (40)$$

If $c \leq 0$, then the claimed property is true because (40) in this case is negative. If $c > 0$, then (40) is positive, equal to zero, and negative, respectively, whenever $\phi(z - s_d)/\Phi(z - s_d)$ is smaller than, equal to, and greater than $s_f c$. By Lemma 2.1 in Baricz (2008), $\phi(\cdot)/\Phi(\cdot)$ is a strictly decreasing function. This implies that $\phi(z - s_d)/\Phi(z - s_d)$ is a strictly increasing function of s_d . Moreover, it holds that $\lim_{s_d \rightarrow 0} \phi(z - s_d)/\Phi(z - s_d) = \phi(z)/\Phi(z)$. Therefore, if $s_f c > \phi(z)/\Phi(z)$, that is, \bar{C} is true, then the threshold \bar{s}_d^V exists and (40) is positive, equal to zero, and negative, respectively, when $s_d \in (0, \bar{s}_d^V)$, $s_d = \bar{s}_d^V$, and $s_d \in (\bar{s}_d^V, \infty)$; if $s_f c \leq \phi(z)/\Phi(z)$, that is, C is true, then (40) is negative for all $s_d > 0$. \square

Proof of Proposition 4 (Extended model). (1) Exploiting condition (20), the objective function of problem (19) can be rewritten as

$$(B - A)Fq - 2A \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E} \left[f_i \left(\frac{q}{I} - d_i \right)^+ \right] - (1 + A) \sum_{i \in \mathcal{I}} \delta^{i-1} \mathbb{E}[f_i d_i]. \quad (41)$$

It is easy to verify that (41) is proper concave in q and tends to $-\infty$ as q tends to ∞ . The first order optimality condition (21), which holds by (20), is thus sufficient and necessary for the optimality of a positive solution to problem (19). If no positive solution satisfies this condition, then zero must optimally solve this problem.

(2) Condition (22) follows from applying Lemma 1 in a manner similar to its application in the proof of Proposition 2. This condition is always satisfied at a positive value of q , as the left hand side of (22) tends to 0 and $I > (1 - B/A)/2$, respectively, when q tends to 0 and ∞ (in the latter case this is due to (20)).

(3) Expression (23) follows from letting q be equal to 0 in (41) and using the properties of the lognormal distribution to express each term $\mathbb{E}[f_i d_i]$. Expression (24) follows from applying Lemma 2 to express each term $\mathbb{E}[f_i (q/I - d_i)^+]$ with $q = q^*$ in (41), exploiting the optimality condition (22), and using (23). Expression (25) follows immediately from (23) and (24). \square

B. The Value of the Demand Forecast Forward Procurement Policy

It follows from (6) and Lemma 2 that the value of the demand forecast procurement policy with the MLN model is

$$V^F = V^S + \{A - B - 2A [\Phi(s_d/2 - cs_f) - K^C \Phi(-(s_d/2 + cs_f))]\} FD. \quad (42)$$

By (41), Lemma 2, and (20), the value of this policy with the extended MLN model is

$$\begin{aligned} V^F &= V^S + (A - B)FD \\ &\quad - 2AD \sum_{i \in \mathcal{I}} \left[\frac{1}{I} \Phi \left(\frac{-\ln(I\beta_i)}{s_{d,i}} + s_{d,i}/2 - c_i s_{f,i} \right) - \beta_i K_i^C \Phi \left(\frac{-\ln(I\beta_i)}{s_{d,i}} - (s_{d,i}/2 + c_i s_{f,i}) \right) \right] \\ &\quad \delta^{i-1} \mathbb{E}[f_i]. \end{aligned} \quad (43)$$

Moreover, if $\beta_i = 1/I$ for all $i \in \mathcal{I}$, then (43) simplifies to

$$V^F = V^S + (A - B)FD - \frac{2AD}{I} \sum_{i \in \mathcal{I}} [\Phi(s_{d,i}/2 - c_i s_{f,i}) - K_i^C \Phi(-(s_{d,i}/2 + c_i s_{f,i}))] \delta^{i-1} \mathbb{E}[f_i]. \quad (44)$$

C. Bounds on the Absolute and Relative Benefits of Optimization

It follows from (9) and (42) that the absolute benefit of optimization under the MLN model is

$$\begin{aligned} V - V^F &= \{ \alpha K^C - A + B + 2A [\Phi(s_d/2 - cs_f) - K^C \Phi(-(s_d/2 + cs_f))] \} FD \\ &= \left\{ -A + B + 2AK^C \left[\Phi(z - s_d) + \frac{\Phi(s_d/2 - cs_f)}{K^C} - \Phi(-(s_d/2 + cs_f)) \right] \right\} FD. \end{aligned} \quad (45)$$

If $c > 0$, it follows from $\Phi(z - s_d) \in (0, 0.5)$, $A > B$, $K^C \equiv \exp(cs_d s_f) > 1$, $\Phi(\cdot) \in (0, 1)$, and (45) that

$$\begin{aligned} V - V^F &< 2AK^C [0.5 + \Phi(s_d/2 - cs_f) - \Phi(-(s_d/2 + cs_f))] FD \\ &< 2AK^C (0.5 + 1) FD \\ &= 3AK^C FD. \end{aligned} \quad (46)$$

If $c = 0$ similar arguments and noticing that $K^C = 1$ imply that

$$V - V^F < 3AK^C FD. \quad (47)$$

If $c < 0$, then $K^C \in (0, 1)$; it thus follows that

$$\begin{aligned}
V - V^F &= \{-A + B + 2AK^C [\Phi(z - s_d) - \Phi(-(s_d/2 + cs_f))] + 2A\Phi(s_d/2 - cs_f)\} FD \\
&< \{-A + B + 2A\Phi(z - s_d) + 2A\Phi(s_d/2 - cs_f)\} FD \\
&< (-A + B + 2A \cdot 0.5 + 2A) FD \\
&= (2A + B)FD.
\end{aligned} \tag{48}$$

Inequalities (46)-(48) imply the validity of (12).

Since $\Phi(z - s_d) \in (0, 0.5)$, it holds that

$$|V| = [1 + A - 2A\Phi(z - s_d)] K^C FD > (1 + A - 2A \cdot 0.5) K^C FD = K^C FD. \tag{49}$$

The inequalities in (13) follow from (12), (49), and $K^C = 1$ if $c = 0$.

References

- Baricz, A. 2008. Mills' ratio: Monotonicity patterns and functional inequalities. *Journal of Mathematical Analysis and Applications* **340**(2) 1362–1370.
- Rockafellar, R. T. 1970. *Convex Analysis*. Princeton University Press, Princeton, NJ, USA.