

---

# VOLATILITY DYNAMICS OF NYMEX NATURAL GAS FUTURES PRICES

---

**HIROAKI SUENAGA\***  
**AARON SMITH**  
**JEFFREY WILLIAMS**

We examine the volatility dynamics of NYMEX natural gas futures prices via the partially overlapping time-series model of Smith (2005. *Journal of Applied Econometrics*, 20, 405–422). We show that volatility exhibits two important features: (1) volatility is greater in the winter than in the summer, and (2) the persistence of price shocks and, hence, the correlations among concurrently traded contracts, displays substantial seasonal and cross-sectional variation in a way consistent with the theory of storage. We demonstrate that, by ignoring the seasonality in the volatility dynamics of natural gas futures prices, previous studies have suggested sub-optimal hedging strategies. © 2008 Wiley Periodicals, Inc. *Jrl Fut Mark* 28:438–463, 2008

The authors thank Harry Bloch and an anonymous referee for their helpful comments.

\*Correspondence author, Hiroaki Suenaga, School of Economics and Finance, Curtin University of Technology, GPO Box U1987, Perth, WA 6845, Australia. Tel: 61-8-9266-4480, Fax: 61-8-9266-3026, E-mail: hiroaki.suenaga@cbs.curtin.edu.au

Received November 2006; Accepted July 2007

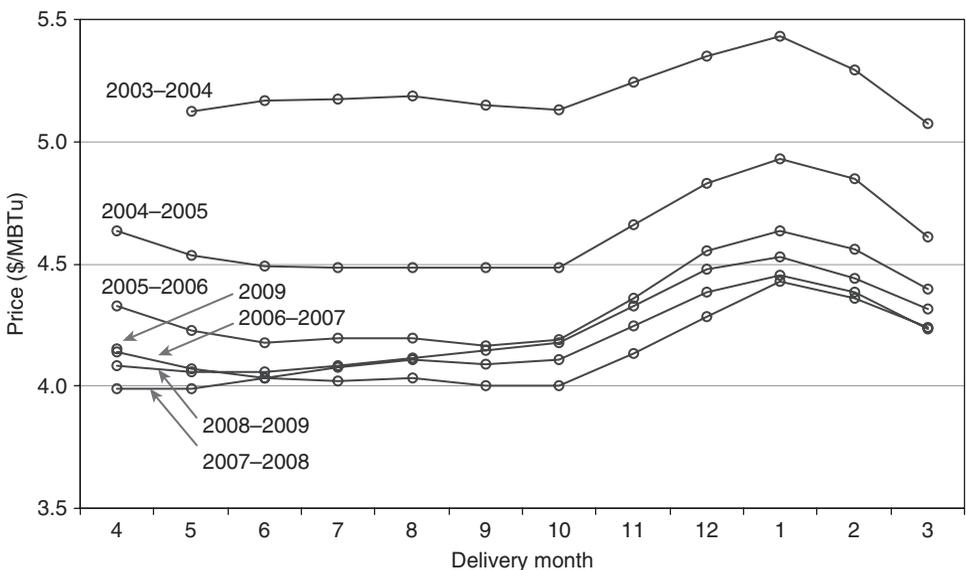
- 
- *Hiroaki Suenaga is a Research Fellow at the School of Economics and Finance, Curtin University of Technology, Perth, Western Australia, Australia.*
  - *Aaron Smith is an Associate Professor at the Department of Agricultural and Resource Economics, University of California, Davis, Davis, California, USA and a Member of the Giannini Foundation, USA.*
  - *Jeffrey Williams is the D. Barton DeLoach Professor, Department of Agricultural and Resource Economics, University of California, Davis, Davis, California, USA and a Member of the Giannini Foundation, USA.*

**INTRODUCTION**

Natural gas demand in the United States peaks between December and March because of residential heating during the winter. Although this unremarkable fact implies higher average prices in winter, intertemporal arbitrage mitigates the price effect; the winter price equals the off-peak (spring and summer) price plus the cost of carry. This pattern is implied by the theory of storage, which asserts that the equilibrium constellation of spot and futures prices represents the prices at which the marginal benefit of current consumption equals the expected marginal benefit of storing a commodity for future consumption (Williams & Wright, 1991). This relationship does not hold between two contracts' prices if inventory is effectively zero at any point between their maturity dates.

The typical seasonality in the constellation of natural gas prices is illustrated in Figure 1, which displays all 72 NYMEX futures prices as of April 1, 2003. Those contracts, although stretching six years into the future, are aligned in Figure 1 to the conventional April–March year in natural gas. This alignment reveals the strong increase in price during the fall and early winter, precisely when stocks are peaking. Unlike for most commodities, storage facilities for natural gas can reach collective capacity during this period, which leads to price relationships that imply significant seasonality in the marginal cost of storage.

Strong seasonality in demand and storage also implies highly nonlinear volatility dynamics. Volatility is naturally high in the winter contracts because of the high marginal cost of natural gas production and the inelastic winter demand



**FIGURE 1**  
 NYMEX natural gas settlement prices as of April 1, 2003.

mean that shocks of even a small magnitude can cause a large price swing. At the same time, the high natural gas inventory intended for all winter months reduces volatility of the early winter contracts, because shocks can be accommodated, albeit partially, by releasing or absorbing inventory. Such flexibility is inevitably lower later in the winter. The seasonal storage pattern also implies that prices of winter contracts start fluctuating as early as the preceding off-peak-demand period, during which information arrives about the future availability of stored gas. In contrast, prices of spring and summer contracts should not exhibit substantial movements until the end of the preceding peak-demand season because little inventory is carried over from winter to spring in a normal year.

Much of the previous research on natural gas, not to mention other energy futures, has not accounted for such complex volatility dynamics resulting from the peculiarities of these commodities. One common approach in previous studies of energy futures prices has been to specify the spot price as a function of underlying state variables following stipulated stochastic processes. The suggested model is then used to derive the valuation formula of futures and other derivative contracts, any difference from the observed futures prices being interpreted as representing a risk premium. Examples of this approach are Schwartz (1997) for the NYMEX crude oil and Lucia and Schwartz (2002) for the Nordpool wholesale electricity market. For natural gas, Manoliu and Tompaidis (2002) apply one- and two-factor models to NYMEX futures prices and report that the estimated models exhibit strong seasonal variation, and that deviations from this seasonal mean revert to zero.

A second common approach has been to examine directly the relationships between the spot and futures prices without stipulating the stochastic processes of the observed price series. Two questions are commonly addressed in such studies: (i) if the futures price is unbiased in forecasting the subsequently realized spot price and (ii) if the spot price and the prices of concurrently traded futures contracts exhibit long-run relationships and, if so, how quickly they revert to the long-run relationship after deviating from it. The first question is usually addressed by examining the statistical significance of the difference between the spot (or nearby futures) and the futures prices observed much earlier. Several studies report that NYMEX natural gas futures prices are downward biased in forecasting the subsequently realized spot price and interpret this bias as representing a risk premium (Modjtahedi & Movassagh, 2005; Movassagh & Modjtahedi, 2005; Walls, 1995). For the second question, Root and Lien (2003) and Lien and Root (1999) use various cointegration methods to examine the long-run relationships between the spot price and the prices of concurrently traded futures contracts. They conclude that these prices share a common stochastic trend and respond to deviations from their long-run relationships in such a way that the difference between the prices converges to

zero. This result implies that simultaneously traded futures prices are in one-to-one relationships and therefore have a basis of zero in the long run.

Improper or incomplete understanding of volatility dynamics can lead such studies to erroneous conclusions about the market risk premium or, more generally, the efficiency of futures markets. Williams and Wright (1991) show that the equilibrium futures prices derived from a dynamic rational-expectations model of a seasonal storable commodity are highly nonlinear and non-smooth, which is to say that they cannot be expressed in a reduced-form. Clearly, specified seasonal price cycles can only approximate the true price dynamics. Strong seasonality in storage implies that the difference between the spot price and prices of concurrently traded futures contract, or the basis, representing the cost of carry, varies across contract delivery months. Previous studies on spot–futures relationships have not incorporated such seasonal variation in the basis.

Aside from these modeling issues, most previous studies utilized only a subset of price data available from organized exchanges where multiple contracts with different maturity dates trade concurrently. One practice is to construct a monthly data series by stacking the futures prices from some particular day of the month. Another practice is to construct a daily series by splicing together the nearby futures contracts. Both of these practices not only discard much information, but they also distort the temporal dynamics of the observed price series because of switching from one delivery month to another.

This study examines the volatility dynamics of NYMEX natural gas futures prices using the partially overlapping time-series (POTS) model of Smith (2005). Unlike conventional models of commodity price dynamics, the POTS model treats the daily prices on a given contract as a single time series. As time proceeds, some contracts reach delivery and cease to exist, and others are born and begin trading. In this sense, a set of contracts constitutes a set of partially overlapping time series. Each of these time series behaves as a martingale process, which allows the model to be applied directly to the first difference of the commodity futures price and therefore to avoid misspecifying the price level dynamics. To account for the volatility dynamics, the POTS model uses nonparametric functions. This flexible specification captures the time-to-maturity effect, seasonal volatility, and other nonlinear volatility dynamics resulting from the peculiarities of natural gas as a commodity.

To illustrate the practical importance of these features of the POTS model, we apply our analysis of the volatility dynamics of natural gas futures prices to the standard “optimal hedging” strategy. We show that seasonal and cross-sectional variations in the degree of persistence of the price shocks, together with the seasonality in the arrival of information, imply that the December and summer (June through August) contracts are more effective than other contracts in minimizing the variance of portfolio returns. This result

is contrasted with the results from previous studies, which ignored seasonal and cross-sectional variations in volatility of natural gas prices and which recommended the use of the adjacent delivery contract for hedging.

The rest of the study is structured as follows. In the second section, the POTS model of Smith (2005) is presented and its benefits and drawbacks relative to conventional models of commodity price dynamics are discussed. The third section describes the natural gas futures price data used for estimating the POTS model and provides the estimation results. The fourth section considers the implications of the estimated volatility dynamics of natural gas futures prices for the optimal hedging strategy. It is shown that taking into account the estimated nonlinear volatility dynamics allows significant reduction in the volatility of portfolio returns as compared with a simple strategy utilizing the adjacent delivery contract. The last section concludes the study.

## POTS MODEL

The POTS model of Smith (2005) is a latent factor model of daily futures price changes. In this model, the change from one day to the next of a particular contract's price is a linear combination of the common latent factors and an idiosyncratic term. The common factors represent the fact that contemporaneous futures prices are tied together by intertemporal arbitrage. In a single-factor setting, the model takes the following form:

$$\Delta F_t = \boldsymbol{\theta}_t \boldsymbol{\varepsilon}_t + \boldsymbol{\lambda}_t \mathbf{u}_t, \quad (1)$$

where  $\Delta F_t$  is an  $n_t \times 1$  vector of daily futures price changes with  $n_t$  representing the number of contracts traded on day  $t$ . It comprises  $\Delta F_{m,t} = F_{m,t} - F_{m,t-1}$  where the two subscripts represent the maturity and trading date, respectively. The two subscripts implicitly define the number of days to maturity as  $d = m - t$ . The scalar  $\boldsymbol{\varepsilon}_t$  is a latent factor with  $E[\boldsymbol{\varepsilon}_t] = 0$  and  $E[\boldsymbol{\varepsilon}_t^2] = 1$  and  $\boldsymbol{\theta}_t$  is an  $n_t \times 1$  vector of factor loadings. The  $n_t \times 1$  vector  $\mathbf{u}_t$  denotes the idiosyncratic error with  $\mathbf{u}_t \sim N(0, I_{n_t})$ , and  $\boldsymbol{\lambda}_t$  signifies an  $n_t \times n_t$  diagonal matrix that determines the standard deviation of the idiosyncratic term.

The components of  $\boldsymbol{\theta}_t$  and  $\boldsymbol{\lambda}_t$  are specified as deterministic functions of maturity date and time to maturity. That is,

$$\begin{aligned} \theta_{m,t} &= \theta(d,m), \\ \lambda_{m,t} &= \lambda(d,m). \end{aligned} \quad (2)$$

For identification, we specify  $E[\boldsymbol{\varepsilon}_t \mathbf{u}_{m,t}] = 0$  for all  $m$  and  $t$ , and  $E[\Delta F_t | \mathfrak{S}^{t-1}] = 0$  where  $\mathfrak{S}^{t-1}$  denotes the information set available at  $t - 1$ . These assumptions

assert that a series of daily futures price changes follows a martingale difference sequence, which implies a zero risk premium.<sup>1</sup>

As in Smith (2005), the conditional variance of  $\varepsilon_t$  is specified by the GARCH(1,1) process,

$$E[\varepsilon_t^2 | \mathfrak{S}^{t-1}] = h_t = \omega + \beta h_{t-1} + \alpha E[\varepsilon_{t-1}^2 | \mathfrak{S}^{t-1}]. \quad (3)$$

Because the unconditional variance of  $\varepsilon_t$  equals unity,  $\omega = 1 - \alpha - \beta$ . The last term in (3) is

$$E[\varepsilon_{t-1}^2 | \mathfrak{S}^{t-1}] = \varepsilon_{t-1|t-1}^2 + P_{t-1|t-1},$$

where  $\varepsilon_{t-1|t-1} = E[\varepsilon_{t-1} | \mathfrak{S}^{t-1}]$  and  $P_{t-1|t-1} = E[(\varepsilon_{t-1} - \varepsilon_{t-1|t-1})^2 | \mathfrak{S}^{t-1}]$ , which are obtained through the Kalman filter (Hamilton, 1994).

The POTS model as defined in (1)–(3) is econometrically similar to factor models of commodity price dynamics, such as Schwartz (1997) and Manoliu and Tompaidis (2002) as applied to natural gas. Among the major differences is that the POTS model views the daily prices on a given contract as a single time series and specifies the dynamics of daily price changes. In contrast, conventional factor models specify the dynamics of daily spot price levels and not price changes. By modeling daily futures price changes, which are a martingale difference sequence, the POTS model avoids specifying seasonal or any other deterministic variations in the underlying spot price level and, hence, is free of approximation bias. This approximation bias could be especially large for commodities that exhibit strong seasonal patterns (Suenaga, 2005).

The POTS model also employs flexible functional forms in specifying the factor loadings ( $\theta_{m,t}$ ) and the standard deviation of the idiosyncratic term ( $\lambda_{m,t}$ ) to capture time-to-maturity effects and seasonal volatility as well as other non-linear volatility dynamics resulting from peculiarities of the commodity. In contrast, conventional factor models specify a stochastic process of the spot price dynamics with a small number of parameters, thereby imposing a particular structure on the factor loadings in the futures pricing equation. Even the most complex of these models, including an increased number of factors and/or specifying dynamics of each factor by more complex stochastic process, still specify a simple variance term for the idiosyncratic errors with at most different variances across contract delivery months.<sup>2</sup>

<sup>1</sup>We also estimated the model allowing for a nonzero mean in the log-price change. The estimate of this mean parameter is very small in value, and our main results are unaffected.

<sup>2</sup>Suenaga (2005) illustrated that this difference in the level of complexity of the stipulated variance structure of the latent factors and idiosyncratic errors has a substantial impact on the other model parameters, especially the risk premium.

## DATA AND ESTIMATION RESULTS

### Data

We estimate the model in (1) through (3) using daily settlement price data from the NYMEX natural gas futures markets. The NYMEX started trading natural gas futures contracts in April, 1990. Contracts traded in this market are defined in monthly blocks, with each contract providing for the delivery of 10,000 million British thermal units of natural gas at the Henry Hub located in Louisiana. The market initially traded contracts up to one year before the first calendar day of the delivery month, but has gradually extended this horizon to six years. Contracts were initially traded until seven business days before the first calendar day of the delivery month, but trading was later extended to three business days before the delivery month.

We analyse the daily change in the logarithm of settlement prices using a sample spanning from January 2, 1991 to December 31, 2003. Because distant delivery contracts are not always actively traded, contracts of more than 12 months to maturity are dropped. Excluding these observations leaves a sample of 40,618 prices among 175 contracts.

### Estimation Results

In applying the POTS model to NYMEX natural gas futures prices, we consider the following three specifications for the factor loadings and the variance of the idiosyncratic error.

$$\psi_{m,t} = \psi^m(d) = a_{\psi,0}^m + a_{\psi,1}^m d + \sum_{j=1}^5 \left( a_{\psi,2j}^m \sin\left(\frac{2\pi jd}{d_{\max}}\right) + a_{\psi,2j+1}^m \cos\left(\frac{2\pi jd}{d_{\max}}\right) \right) \quad (4a)$$

$$\begin{aligned} \psi_{m,t} = \psi(t, d) = & a_{\psi,0} + a_{\psi,1} d + \sum_{j=1}^5 \left( a_{\psi,2j} \sin\left(\frac{2\pi jd}{d_{\max}}\right) + a_{\psi,2j+1} \cos\left(\frac{2\pi jd}{d_{\max}}\right) \right) \quad (4b) \\ & + \sum_{j=1}^5 \left( b_{\psi,2j-1} \sin\left(\frac{2\pi j DC(t)}{365}\right) + b_{\psi,2j} \cos\left(\frac{2\pi j DC(t)}{365}\right) \right), \end{aligned}$$

$$\psi_{m,t} = \psi(d) = a_{\psi,0} + a_{\psi,1} d + \sum_{j=1}^5 \left( a_{\psi,2j} \sin\left(\frac{2\pi jd}{d_{\max}}\right) + a_{\psi,2j+1} \cos\left(\frac{2\pi jd}{d_{\max}}\right) \right), \quad (4c)$$

where  $d_{\max} = 365$ ,  $\psi_{m,t} = \lambda_{m,t}$ ,  $\theta_{m,t}$ , and  $DC(t)$  is the day in year of  $t$ . Equation (4a), the most flexible of the three specifications, specifies the factor loadings and the variance of idiosyncratic error as a function of days to delivery of the contract and allows these functions to differ among the 12 delivery months ( $m = 1, \dots, 12$ ). Equation (4c) employs the same functional form as (4a) but does

not allow the function to vary across contracts. Equation (4b) is the intermediate specification. It allows time-to-maturity effect and seasonal variation (terms in the first and second summation operator, respectively), yet restricts these two effects to be identical to all contracts.

In all three specifications, the function becomes more flexible as the number of trigonometric terms,  $j$ , increases. Although this extra flexibility allows the model to fit the observed data better, it also makes the coefficient estimates more sensitive to extreme observations. Using only five trigonometric terms allows sufficient flexibility to capture seasonality and time-to-delivery effects while avoiding excess sensitivity to outliers.

The model in these three specifications is estimated by the method of maximum likelihood with initial values obtained by the iterative approximate-EM method of Smith (2005).<sup>3</sup> In Table I, the results from a specification test shows that specification (4a) is supported over the two more restrictive specifications, and specification (4b) is supported over (4c) by the Schwartz information criterion, calculated as  $SIC_i = k_i \ln(N)/N - 2 \ln(L_i)/N$ , where  $\ln(L_i)$  and  $k_i$  are, respectively, the log-likelihood and the number of parameters in the specification  $i$ , and  $N$  is the sample size. That is, not only does the volatility of daily changes of natural gas futures prices exhibit significant seasonal variation in addition to the time-to-maturity effect, but also these volatility dynamics vary across the 12 delivery months. Given this result, our subsequent discussion focuses only on the volatility dynamics implied by the estimated POTS model in specification (4a).

Table II summarizes the coefficient estimates of the GARCH parameters. In the table, the coefficient estimates of  $\alpha$  and  $\beta$  are 0.092 and 0.898. The sum of the two coefficients, 0.990, is statistically significantly less than unity, indicating

**TABLE I**  
Results of Model Specification Tests

<i>Specification</i>	<i>(4a)</i>	<i>(4b)</i>	<i>(4c)</i>
Description	Mean reversion, seasonality, composite	Mean reversion, seasonality, additive	Mean reversion only
Sample size	40,618	40,618	40,618
Number of parameters	290	48	26
Log-likelihood	151,682.6	148,377.8	147,705.1
SIC	-7.3930	-7.2961	-7.2671

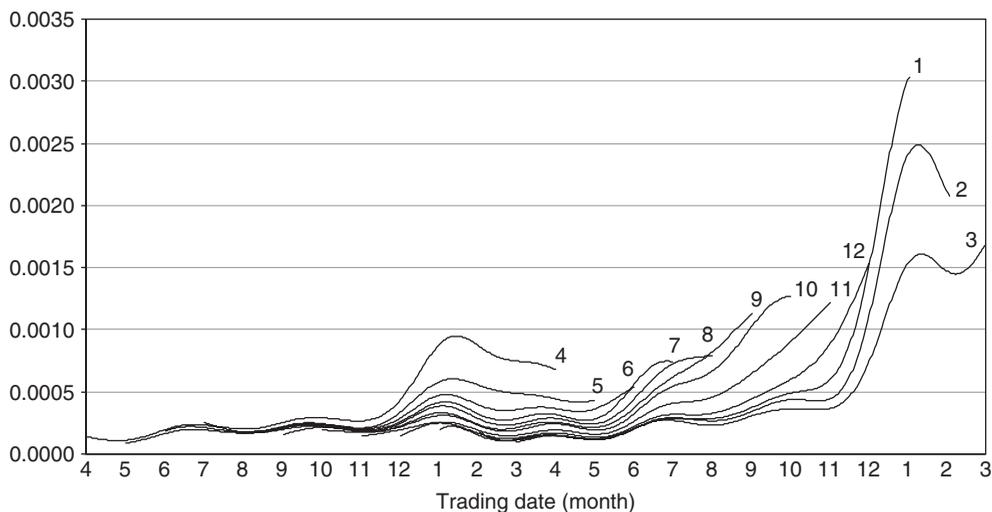
SIC, Schwartz information criterion

<sup>3</sup>This approximate-EM method involves iteration of the following three steps: (1) obtain the predicted values of the latent factor and GARCH conditional variance through the Kalman Filter, (2) maximize the expected complete-data likelihood with respect to the spline parameters, conditional on the predicted values of latent factor and GARCH conditional variance from the first step, and (3) estimate the GARCH parameters holding the spline parameters at the values from step 2.

**TABLE II**  
Maximum Likelihood Estimates of GARCH Parameters

GARCH parameters	Coefficient	Std. err.	t-Ratio
$\alpha$	0.0923	0.0100	9.2005
$\beta$	0.8981	0.0112	80.2181
$\alpha + \beta$	0.9903	0.0040	2.4426*

\*t-Ratio for  $\alpha + \beta$  is with respect to 1.



**FIGURE 2**  
Variance of daily log-price changes implied by the estimated POTS model.

that the conditional volatility of the common underlying factor is highly persistent yet is stationary.

Figures 2 and 3 plot the unconditional variance of daily price changes of each of 12 contracts, which is computed as  $\theta_{m,m-d}^2 + \lambda_{m,m-d}^2$  for  $d$  ranging from 0 to 365 days to maturity, using the estimated factor loadings,  $\theta_{m,t}$ , and the standard deviation of the idiosyncratic term,  $\lambda_{m,t}$ . These estimates of unconditional variance are aligned with the trading day in Figure 2 and with the number of days to delivery in Figure 3. In Figure 3, the associated 95% confidence intervals are also presented.<sup>4</sup> The two figures exhibit at least four interesting

<sup>4</sup>The standard error estimates of the unconditional variance of daily log-price changes were obtained by the delta method. In Figure 3, the estimated unconditional variance is significantly positive for all 12 delivery months and for all trading dates within one year before maturity. The confidence intervals tend to be wide as the contract approaches its maturity date, particularly for the January through March contracts. The fairly large confidence intervals obscure the significance of the difference among the 12 delivery months in their estimated variance of daily price changes, however. The specification test in Table I reveals the differences to be substantial overall.

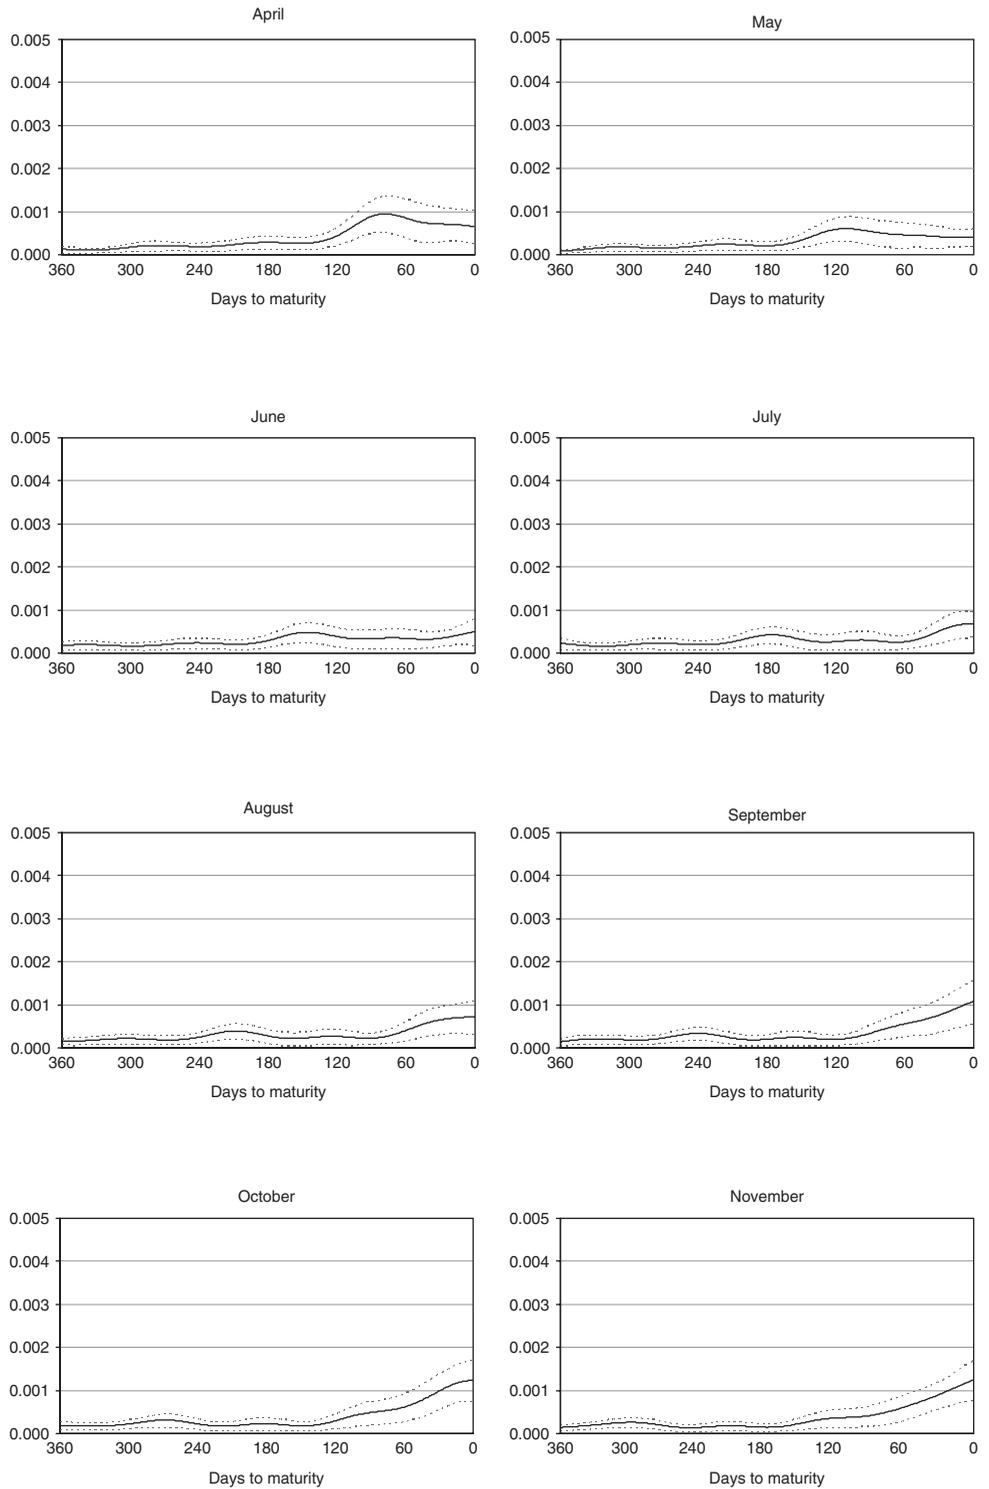
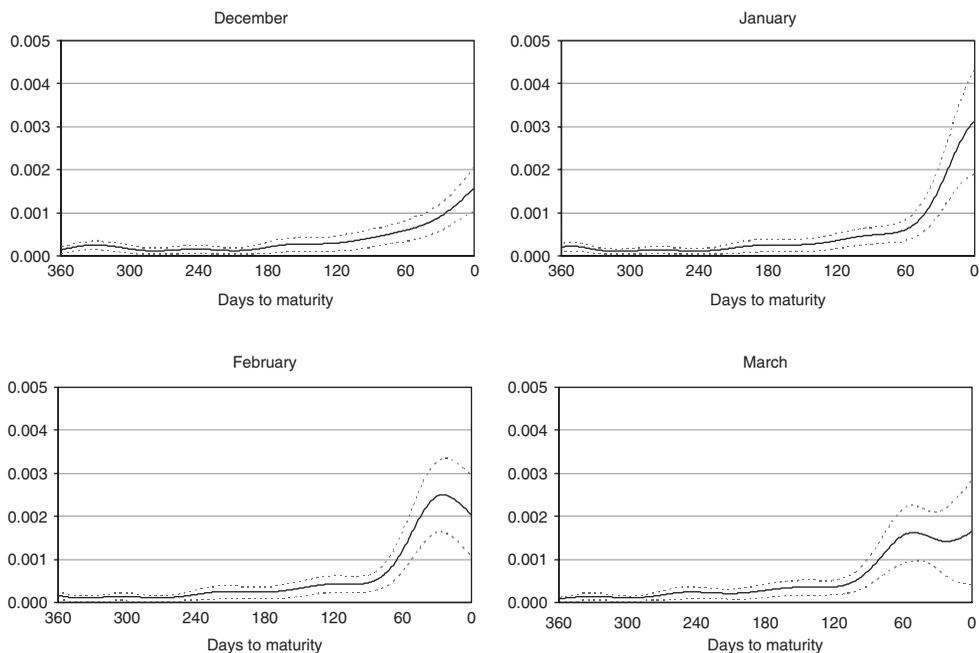


FIGURE 3 (Continued)



**FIGURE 3**

Variance of daily log-price changes implied by the estimated POTS model with the associated confidence intervals.

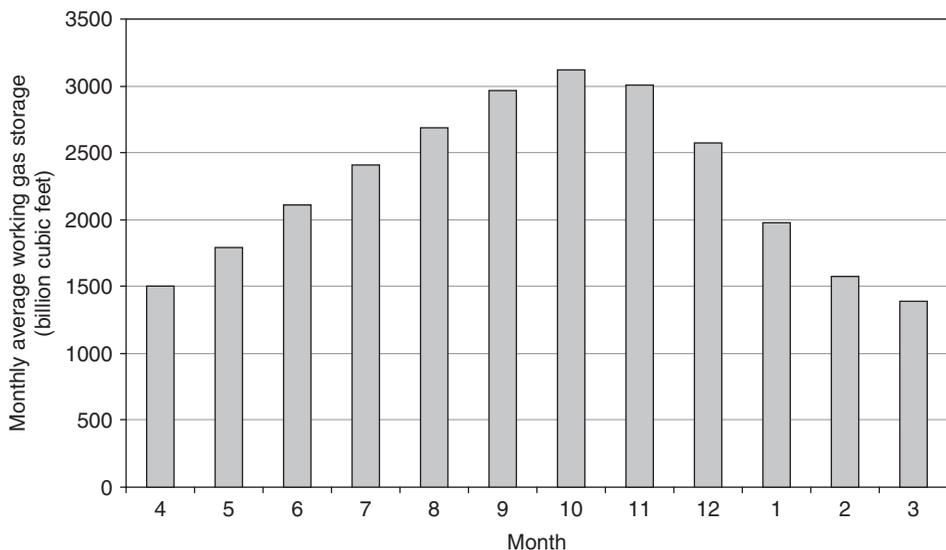
features. First, on any given date, the volatility of contracts that are closer to maturity exceeds that of more distant contracts. This feature, often called the Samuelson (1965) effect, indicates that shocks to spot prices are expected to dissipate somewhat over time. Second, the volatility of close-to-maturity contracts is substantially higher for winter contracts than for late spring or summer contracts. Third, during the period from early May to late September, volatility increases for all contracts maturing before the end of the following peak season. Finally, during the early winter (early November to mid-January), volatility rises for all 12 contracts, although the January to April contracts display much larger increases than the other contracts.

The last three of the four features observed in Figure 2, those pertaining to the seasonal variations in the volatility of daily log-price changes of natural gas futures contracts, follow naturally from seasonality in the demand for natural gas. Nearby winter contracts exhibit high price volatility because demand reaches contemporaneous supply capacity in winter. Even though high winter inventory allows price shocks of some magnitude to be absorbed, the high winter volatility depicted in Figure 2 indicates that such a price buffering is only limited. Summer contracts exhibit low nearby volatility because low demand relative to supply means that demand shocks can be absorbed by altering the amount injected into storage.

The gradual volatility increase in the late spring and summer (May to late September) reflects the arrival of important market information. During this period, inventory continuously changes as excess natural gas production is stored for use in the subsequent peak-demand season. Gas stored in this period will not be carried over to the following off-peak-demand season unless gas demand in the coming peak-demand season is unusually low. Such information will not be revealed until actual demand and/or supply conditions are realized. Hence, price volatility increases from May to late September, but only for the contracts deliverable during the coming peak-demand season and earlier.

During early winter, new information provides strong signals about the imminent peak season as well as the likely amount of inventory carryover from the current peak-demand season to subsequent months. Nonetheless, the amount of gas carried over from the peak to the subsequent off-peak season is very small relative to the amount stored over the off-peak season in a normal year. Thus, nearby contracts exhibit high volatility during this early winter period, but volatility also increases, albeit only marginally, for contracts that mature in the summer. By mid-January, most of the information about the current peak season has arrived and price volatility begins to decline.

These seasonal patterns in price variance correspond to the seasonal patterns in gas storage. Temporal arbitrage induces a large amount of natural gas to be carried over from the off-peak-demand season (April through June) to the peak-demand season (December to March) so that the price difference between these seasons is on average equal to the cost of carry. Figure 4 illustrates this storage



**FIGURE 4**

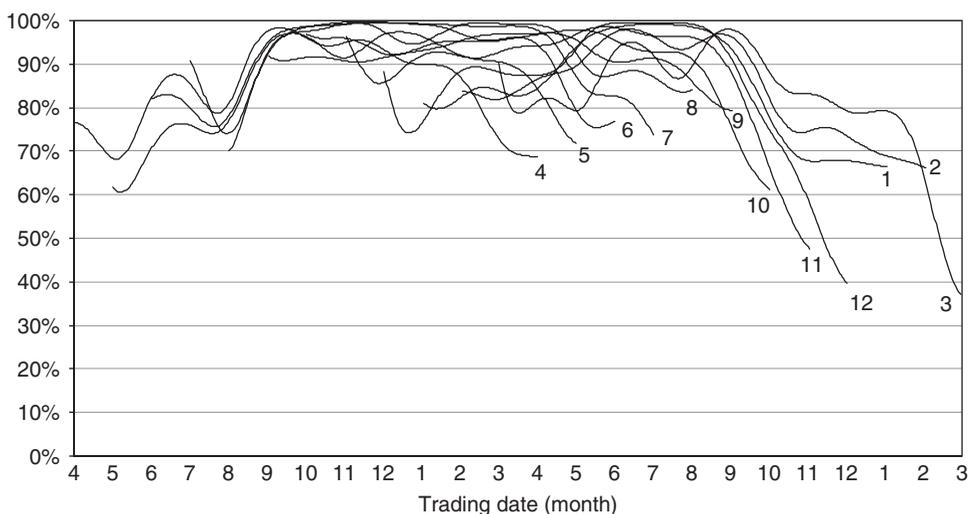
U.S. natural gas underground storage—monthly average of working gas for 1976–2005.

pattern. In the figure, U.S. nationwide working gas storage is the lowest in March, after which it gradually increases throughout spring and summer until it reaches an annual high in October or November. A rapid decrease of inventory from November through the following February indicates that a large amount of stored gas is withdrawn to meet with a high demand for heating energy.

Table III summarizes the proportion of the variance explained by the common factor,  $\epsilon_t$ . It indicates that the model explains, on average, about 83.5% of daily price variation. In general, spring and summer contracts are more closely related to the common factor than winter contracts. Figure 5 illustrates, for

**TABLE III**  
Proportion of the Variance Explained by a  
Common Factor

Overall	0.8348
By contract	
1	0.7587
2	0.7528
3	0.7498
4	0.8538
5	0.8897
6	0.9217
7	0.9190
8	0.9225
9	0.9025
10	0.8640
11	0.8326
12	0.7631



**FIGURE 5**  
Share of price variation explained by the common factor.

each of the 12 maturities, how the proportion of total price variance explained by the common factor changes over the time ahead of expiration. We measure this proportion by  $\theta_{m,m-d}^2 / (\theta_{m,m-d}^2 + \lambda_{m,m-d}^2)$ . For all contracts, this proportion drops substantially below 100% in the last few months before maturity. This pattern indicates that, as delivery approaches, the futures price increasingly reflects local conditions at the Henry Hub and less conditions in natural gas markets in the rest of North America. Such a pattern is not surprising for natural gas. Because the futures contract for natural gas represents a month-long flow through the Henry Hub, rather than the more traditional warehouse receipt for the commodity in store, local congestion in pipelines can disconnect nearby futures from markets elsewhere. (For the same reasons of network congestion, those local markets elsewhere sometimes have spot or nearby forward prices either at steep discounts or steep premiums to the average throughout the network.)

Figure 5 also shows that, for contracts that deliver between October and March, the share explained by the common factor drops substantially around the middle to the end of August and keeps decreasing afterward. That is, much of the post-August price variation in these contracts emanates from information of a short-term nature that does not affect the amount of carryover to the following off-peak season. Coupled with the high volatility of these contracts in this period (see Fig. 2), this observation implies that the high inventory accumulated by late summer allows only limited buffering of current market shocks. This lack of ability to absorb shocks likely arises from the high marginal cost of inventory adjustment, which increases with the inventory level because injecting gas back into underground storage requires the gas to be at ever higher pressures.

In this same post-August period, contracts that deliver after the end of the following peak season exhibit low total volatility (see Fig. 2), but a high proportion of their variation is captured by the common factor (see Fig. 5). Thus, the common factor represents information about market conditions in the following off-peak season, but relatively little such information arrives in this period. The fact that the October to March contracts are weakly related to the factor in this period is consistent with the small inventory carryover from the peak season to the following off-peak season. This lack of carryover reduces the potential for arbitrage across these two periods, and therefore it weakens the link between the prices of contracts that deliver before the end of the peak season and those that deliver after the end of the peak season.

In sum, the estimated POTS model reveals that the volatility of natural gas futures prices exhibits not only the usual Samuelson effect but also strong seasonal and cross-sectional variation. Volatility is higher for winter contracts than for other contracts. Volatility also increases from early May to September

for all contracts maturing by the end of the following peak-demand season and from early November through mid-January for the January to April contracts. The correlation between the daily price changes of concurrently traded contracts tends to be highest during these two periods. Overall, our model illuminates the complex dynamics of natural gas price volatility, which previous studies cannot discern because their econometric models presume a volatility specification that is too simple.

## OPTIMAL HEDGING STRATEGY

This section extends our analysis of the volatility dynamics of natural gas futures prices to investigate the implications for hedging. A simple hedging strategy is considered, in which a hedger has a spot position,  $Q$ , at time  $t$  and simultaneously takes a short position in  $X$  futures contracts for delivery at  $\tau > t$ . At  $t + k < \tau$ , the hedger clears its position by selling  $Q$  units in the spot market and buying  $X$  futures contracts for delivery at  $\tau$ .

The hedger's change in wealth from  $t$  to  $t+k$ , ignoring the interest rate, is

$$W_{t+k} = (S_{t+k} - S_t)Q - (F_{t+k,\tau} - F_{t,\tau})X = (\Delta S_{t+k} - \eta \Delta F_{t+k,\tau})Q, \quad (5)$$

where  $S_i$  is the spot price at  $i$ ,  $F_{i,\tau}$  is the period  $i$  price of the futures contract for delivery at  $\tau$ , and  $\eta = X/Q$  is the hedge ratio. The variance of  $W_{t+k}$  is

$$V[W_{t+k}] = (V[\Delta S_{t+k}] + \eta^2 V[\Delta F_{t+k,\tau}] - 2\eta \text{cov}[\Delta S_{t+k}, \Delta F_{t+k,\tau}])Q^2, \quad (6)$$

which is minimized by

$$\eta_{t,\tau}^* = \text{cov}[\Delta S_{t+k}, \Delta F_{t+k,\tau}] / V[\Delta F_{t+k,\tau}]. \quad (7)$$

Substituting (7) into (6) yields the minimized variance

$$V[W_{t+k} | \eta_{t,\tau}^*] = V[\Delta S_{t+k}] (1 - \rho_{t+k,\tau}^2) Q^2, \quad (8)$$

where  $\rho_{t+k,\tau}$  is the correlation between spot and futures contract for delivery  $\tau$  in their price changes over period  $t$  to  $t+k$ .

Alternatively, for a hedger minimizing the variance of the portfolio return,  $r = r_S - \eta^r r_F$ , the optimal hedge ratio and the associated minimum variance are

$$\eta_{t,\tau}^{r*} = \text{cov}[\Delta \ln S_{t+k}, \Delta \ln F_{t+k,\tau}] / V[\Delta \ln F_{t+k,\tau}], \quad (9)$$

$$V[r_{t+k} | \eta_{t,\tau}^{r*}] = V[\Delta \ln S_{t+k}] (1 - \rho_{t+k,\tau}^{r*2}) Q^2, \quad (10)$$

where  $\rho_{t+k,t}^r$  represents the correlation between the log spot price and log price of the futures contract for delivery at  $\tau$  over the period  $t$  to  $t + k$ .

Three remarks should be made about these hedging strategies. First, many organized exchanges concurrently trade multiple contracts with different maturity dates. Thus, although the futures contract to be included into the portfolio,  $\tau$ , is exogenous to the above strategy, it should be a decision variable for the hedger. Expressions (8) and (10) indicate that given the time of entry,  $t$ , and hedging horizon,  $k$ , the hedger should include in its portfolio the futures contract for which the price change has the highest correlation with the spot price change. In practice, only a finite number of contracts are traded on any given day and one needs to calculate the optimal hedge ratio only for the futures contract with the highest correlation to the spot price over the hedging horizon  $k$ .

Second, the time of entry,  $t$ , and hedging horizon,  $k$ , should also be endogenous to the hedger's decision. The choice of  $t$  and  $k$  is important particularly for a commodity with strong seasonality in mean price. Given a priori knowledge of such a seasonal pattern, a hedger should not hold a spot position from the peak to the off-peak-demand season, for such a practice would yield, at least on average, a negative return. The hedge ratios in (7) and (9) are optimal only given that the decision to carry over from  $t$  to  $t + k$  is predetermined. In practice, a position should be held only if the expected wealth,  $E[W_{t+k}]$ , exceeds the value associated with the minimum variance of portfolio return.

Finally, expressions (7)–(10) indicate that both the optimal hedge ratio and the futures contract included into the portfolio depend on two attributes: (i) the covariance of spot and futures price and (ii) the variance of futures price changes. Specifications on the volatility dynamics of spot and futures prices play key roles in determining empirical estimates of these statistics and, hence, the optimal hedging strategy. The POTS model, by allowing seasonal and cross-sectional variations in the factor loading and the variance of the idiosyncratic error, yields an optimal hedge ratio that varies by  $\tau$ ,  $t$ , and  $k$ . In contrast, conventional models of commodity price dynamics determine factor loadings by the contract's time to maturity and a small number of parameters defining the stochastic processes of the underlying factors. Owing to this restrictive specification, the optimal hedge ratio implied by these models does not vary by contract delivery date. In particular, the simple one-factor mean-reversion model considered by Schwartz (1997) and the two-factor model by Manoliu and Tompaidis (2002) both imply that the optimal portfolio always includes the nearby contract. In other words, the specification choice for the stochastic processes of the underlying factors determines the optimal hedging strategy. The standard regression models considered for the analysis of the spot–futures price relationship are similarly incapable of implying cross-sectional and

seasonal variations in the optimal hedge ratio owing to a simple variance structure assumed for the disturbance term.

### Optimal Hedging Strategy Implied by the POTS Model

We evaluate the optimal hedge ratio based on the unconditional variance of daily price changes implied by the estimated POTS model.<sup>5</sup> In doing so, the nearby futures price is used as a proxy for the spot price. As noted in Estimation Results section, as delivery approaches, the futures price increasingly reflects conditions at the Henry Hub. In this sense, the nearby contract approximates a spot price against which market participants may want to hedge. We consider the case in which  $t$  and  $k$  are predetermined and the hedger chooses which futures contract to include in the portfolio,  $\tau$ , and the hedge ratio. We derive the optimal solution for two cases with different holding periods: (i) a hedger who carries a short position only for a single day ( $k = 1$ ) and (ii) a hedger who carries a short position for one month, starting at the first day of each calendar month and ending at the last day of the same month. These two cases are considered to determine whether very short holding periods are less sensitive to the modeling of price volatility.

From the POTS model, the unconditional variance of daily futures price changes is given as

$$E[\Delta F_t \Delta F_t'] = \theta_t \theta_t' + \lambda_t \lambda_t'. \quad (11)$$

With the assumption  $\mathbf{u}_t \sim N(0, I_{n_t})$ , the diagonal and off-diagonal elements are

$$E[\Delta F_{m,t}^2] = \theta(d, m)^2 + \lambda(d, m)^2, \quad (12)$$

$$E[\Delta F_{m,t} \Delta F_{m',t}] = \theta(d, m)\theta(d', m') \text{ for } m, m' > t \text{ and } m \neq m',$$

where  $d = m - t$  and  $d' = m' - t$ . Using expressions (12), the correlation between daily price changes of two futures contracts for delivery  $m_1$  and  $m_2$  is

$$\rho_{m_1, m_2, t} = \frac{\theta(d_1, m_1)\theta(d_2, m_2)}{(\theta(d_1, m_1)^2 + \lambda(d_1, m_1)^2)^{1/2}(\theta(d_2, m_2)^2 + \lambda(d_2, m_2)^2)^{1/2}} = \pi_{m_1, t} \pi_{m_2, t}, \quad (13)$$

<sup>5</sup>One can also evaluate the optimal hedge ratio using the model's implied *conditional* variance. Unlike the *unconditional* variance used in the main text, the *conditional* variance depends on the historical movements of the natural gas futures price. We only present the evaluations based on the *unconditional* variance because our objective is to emphasize the need to model seasonality.

where  $d_i = m_i - t$  and  $\pi_{m_i,t} = \theta(d_i, m_i)(\theta(d_i, m_i)^2 + \lambda(d_i, m_i)^2)^{-1/2}$  is the square root of the share of the total variance of contract  $i$  explained by the common factor. Equation (13) indicates that the variance of the portfolio return is minimized when the portfolio includes the futures contract for which the largest share of price change is accounted for by the common factor.

For a day-long holding period, the optimal hedge ratio and the minimized variance are

$$\eta_{m_1,m_2,t}^* = \frac{\theta(d_1, m_1)\theta(d_2, m_2)}{\theta(d_2, m_2)^2 + \lambda(d_2, m_2)^2}, \tag{14}$$

$$V[W_{t+1}|\eta_{m_1,m_2,t}^*] = (\theta(d_1, m_1)^2 + \lambda(d_1, m_1)^2)(1-\rho_{m_1,m_2,t}^2)Q^2. \tag{15}$$

For the hedging strategy with the month-long holding period over  $t_1$  to  $t_2$ , we need the expressions for the unconditional variance of daily futures price changes, which, with the martingale property assumed in the model, is simply the sum of the daily price changes over this period:

$$E[\Delta F_{t_1,t_2}\Delta F'_{t_1,t_2}] = \sum_{s=t_1}^{t_2} (\boldsymbol{\theta}_s\boldsymbol{\theta}'_s + \boldsymbol{\lambda}_s\boldsymbol{\lambda}'_s) \tag{16}$$

with its elements,

$$E[\Delta F_{m,t_1,t_2}^2] = \sum_{s=t_1}^{t_2} \theta(m - s, m)^2 + \lambda(m - s, m)^2, \tag{17}$$

$$E[\Delta F_{m,t_1,t_2}\Delta F_{m',t_1,t_2}] = \sum_{s=t_1}^{t_2} \theta(m - s, m)\theta(m' - s, m')$$

for  $m, m' > t_2, m \neq m'$ .

The correlation between the two contracts over the horizon  $t_1$  to  $t_2$  is

$$\rho_{m_1,m_2,t_1,t_2} = \frac{\sum_{s=t_1}^{t_2} \theta(m_1 - s, m_1)\theta(m_2 - s, m_2)}{(\sum_{s=t_1}^{t_2} \theta(m_1 - s, m_1)^2 + \lambda(m_1 - s, m_1)^2)^{1/2}(\sum_{s=t_1}^{t_2} \theta(m_2 - s, m_2)^2 + \lambda(m_2 - s, m_2)^2)^{1/2}}. \tag{18}$$

The optimal hedge ratio and the associated minimum variance are

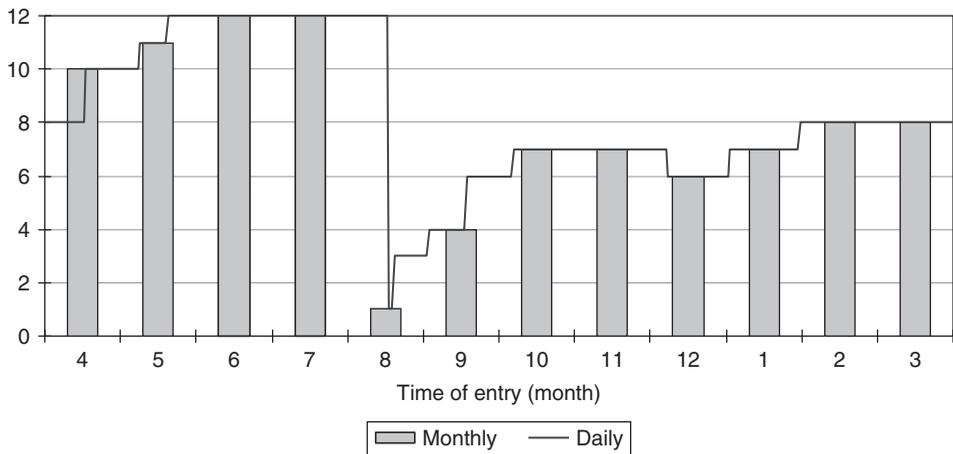
$$\eta_{m_1, m_2, t_1, t_2}^* = \frac{\sum_{s=t_1}^{t_2} \theta(m_1 - s, m_1) \theta(m_2 - s, m_2)}{\sum_{s=t_1}^{t_2} \theta(m_2 - s, m_2)^2 + \lambda(m_2 - s, m_2)^2}, \tag{19}$$

$$V[W_{m_1, m_2, t_1, t_2} | \eta_{m_1, m_2, t_1, t_2}^*] = \sum_{s=t_1}^{t_2} (\theta(m_1 - s, m_1)^2 + \lambda(m_1 - s, m_1)^2) (1 - \rho_{m_1, m_2, t_1, t_2}^2) Q^2. \tag{20}$$

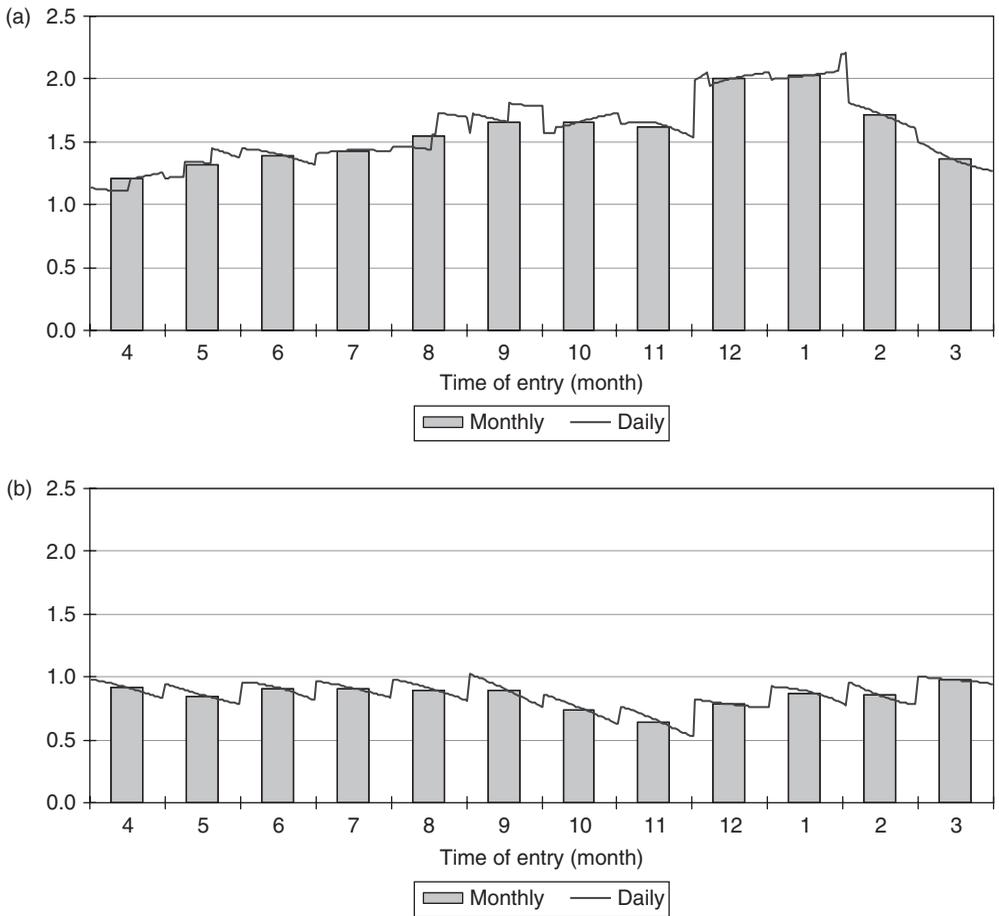
Equations (14) and (19) indicate that the optimal hedge ratio is a function of the time of entry,  $t$ , the hedging horizon,  $k$ , and the delivery period of the futures contract in the portfolio,  $m_2$ . In particular, the optimal hedge ratio increases with the variance of the nearby futures price that is attributable to the common factor and/or the proportion of the futures price variance explained by the common factor. In other words, a hedger should take a large short position when the nearby futures price is very volatile and it is strongly related to the futures prices of subsequent delivery. Even when the nearby futures price is very volatile, a hedger’s position is small if its price movement is not closely related to the other concurrently traded contracts.

### Optimal Hedging Strategy for Natural Gas

Figures 6–9 illustrate the futures contract included into the optimal portfolio, the optimal hedge ratio, and the minimum variance attained by the optimal portfolio for each of the two hedging horizons. First, for a daily hedging strategy,

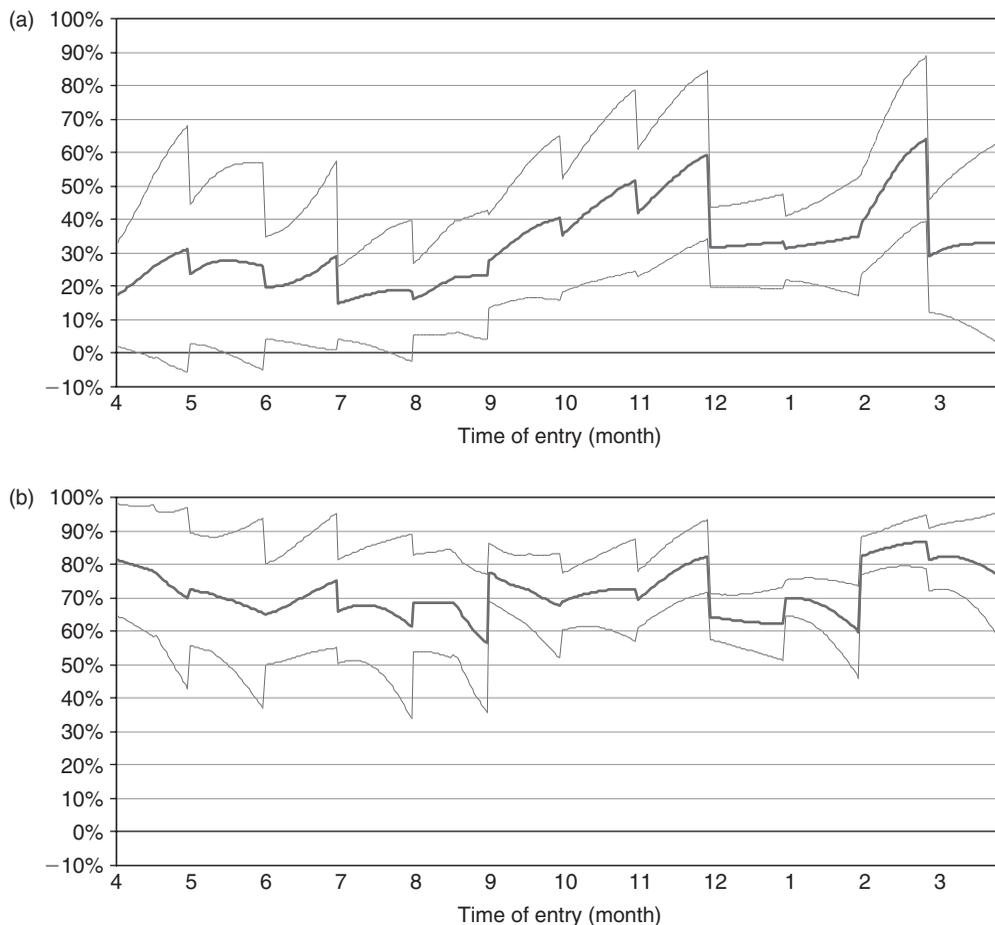


**FIGURE 6**  
Delivery month of the futures contract included in the optimal portfolio.



**FIGURE 7**  
 Optimal hedge ratio: (a) portfolio suggested by the estimated POTS model and (b) portfolio including the second-position contract.

the optimal portfolio frequently includes four contracts: the December contracts for the period between mid-May and mid-August and either the June, July, or August contract for the period between mid-September and mid-April in the following year. These four contracts are often the ones used because they exhibit the highest share of their price volatility explained by the common factor in the relevant period. Other contracts are rarely included in the portfolio. Because short-dated contracts exhibit substantial idiosyncratic volatility, the optimal portfolio never includes these contracts. The three-month-ahead contract is the shortest horizon and is used only in the first half of April and the first half of August where the variance minimizing contract switches gradually from the August to winter and from the December to summer contracts, respectively.

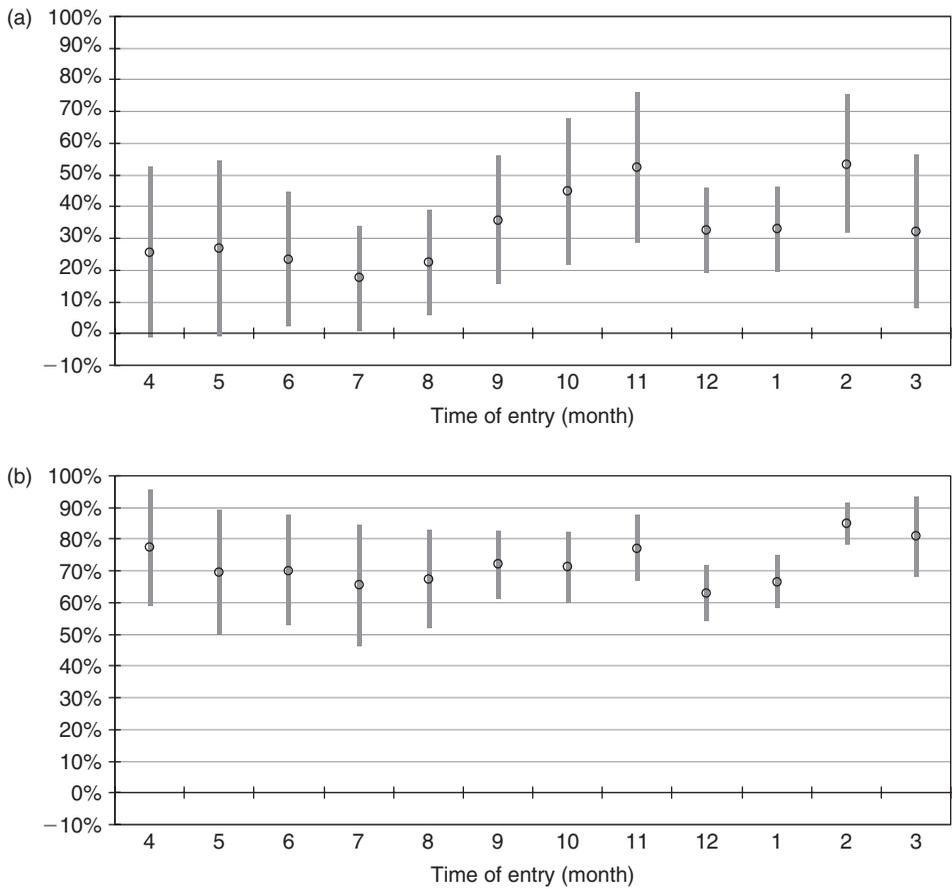


**FIGURE 8**

Minimum variance attained by the optimal portfolio—daily hedging horizon: (a) relative to the variance of unhedged portfolio and (b) relative to the variance of portfolio including the second-position contract.

The predominance of the June, July, August, and December contracts is consistent with our previous discussion relating seasonality of price volatility to that of natural gas demand and storage. From the beginning of the calendar year, the optimal portfolio includes the July contract and then it switches to the August contract. Contracts for earlier maturity are not used owing to their high idiosyncratic volatility. This high idiosyncratic volatility represents the Samuelson effect for the March contract, whereas it stems from low gas storage for the April to June contracts—price movements of these contracts in response to demand, supply, and other market shocks are not linked to one another through available storage.

In the middle of April, the optimal contract for hedging purposes switches from the August to distant contracts as the August contract becomes subject to



**FIGURE 9**

Minimum variance attained by the optimal portfolio—monthly hedging horizon: (a) relative to the variance of unhedged portfolio and (b) relative to the variance of portfolio including the second-position contract.

the maturity effect. This transition is rather quick with the August contract replaced by the December contract by mid-May. The September through November contracts are used only for a short duration because of their high idiosyncratic variance. The next transition takes place in the middle of August when the optimal portfolio switches to the March contract as idiosyncratic variance of the December contract increases rapidly. The January and February contracts are not used in this transition, because a large share of their price variation is contract specific. The March and April contracts are optimal only for a short period because price movements of these two contracts are contract specific owing to the low inventory during those two months.

Figure 7a plots the optimal hedge ratio as a function of the date within the gas year. As shown in (14), the optimal hedge ratio is simply the ratio of the

covariance between the nearby and the futures contract to the variance of the futures contract included in the portfolio. In the figure, the optimal hedge ratio is always above one, indicating that the covariance always exceeds the variance of the futures contract. This is because the futures contracts included into the optimal portfolio are at least three months away from delivery and, hence, display little variation, a dominant share of which reflects the common factor. (Whether a hedge ratio above unity makes sense is a question for the whole theory of “optimal hedging”.) The optimal hedge ratio is particularly high in December and January, during which the price variance of the nearby contract is largely attributable to the common factor, yielding large covariance. Even though the nearby futures contract exhibits high price volatility from September to November, the optimal hedge ratio is only moderate in these months because much of the price variation is contract specific and, hence, has low covariance.

Figure 7b plots the optimal hedge ratio for the portfolio including the second-position contract, which is the optimal solution according to the one- and two-factor models considered by Manoliu and Tompaidis (2002).<sup>6</sup> The mean-reversion process assumed for the underlying factor in their model assures that the factor loading increases monotonically as the contract approaches maturity. In addition, they specify the variance of the idiosyncratic error to be constant over the entire horizon. Consequently, the price correlation must always be the highest for the two contracts with the minimum distance in their maturity dates.

Two observations become clear from the comparison of the two hedging strategies in Figure 7b. First, the optimal hedge ratio is substantially lower for the portfolio including the second-position contract than for the optimal portfolio implied by the POTS model. This is simply because the share of the price variance explained by the common factor is smaller for the second-position than for more distant contracts. Second, the optimal hedge ratio decreases toward the end of every month. This is, again, because the share of the price variance explained by the common factor decreases owing to the Samuelson effect.

Figure 8a illustrates the variance of the optimal portfolio relative to the variance of the unhedged portfolio, i.e., the total variance of the nearby futures price. The figure shows that the optimal hedging strategy implied by the estimated POTS model reduces price risk substantially and significantly. The variance is reduced to less than 40% of the variance of the nearby futures, except that it is above 40% in October, November, and February. The magnitude of variance reduction is small in these three months as the nearby futures prices (the November, December, and March contracts) are inherently very volatile

<sup>6</sup>See Tables 4 and 5 on pages 38 and 39 of Manoliu and Tompaidis (2002). Numbers in Figure 7b were calculated by using the volatility estimates of the POTS model in (14).

and their price volatility is contract specific. Quite noticeably, the variance of the optimal portfolio is below 35% of the variance of the nearby futures price for December and January during which the nearby futures (January and February contract) prices are very volatile. For each of the 12 months, the minimum variance attained by the optimal portfolio is higher toward the end of each month, simply because the Samuelson effect raises the contract-specific volatility of the nearby futures. Nonetheless, the confidence interval of the ratio of the two variances is below one throughout the year.

Figure 8b compares the optimal hedging strategy implied by the POTS model with the straightforward strategy of utilizing the second-position contract. The figure indicates that the strategy based on the POTS model attains the portfolio variance that is about 15–45% below the variance of the portfolio utilizing the second-position contract. This is because the futures contract included in the former strategy is less subject to contract-specific variations than in the second-position contract. The percentage reduction obtained in the portfolio variance is statistically significant—the confidence interval of the ratio of the two variances is below one throughout the year.

The optimal hedging strategies for a monthly horizon yield essentially the same results as those for a daily horizon. The optimal hedge ratio and the minimum variance attained by the optimal contract (Figs. 7 and 9, respectively) are almost identical to the monthly averages of their corresponding values for daily hedging horizon—price risk is reduced to half the size of price variance of the nearby contract and the hedge ratio ranges between 1.2 and slightly above 2.0, owing to a small variance of the futures contract included in the portfolio. Three summer contracts (June, July, and August) and the December contract are the most commonly included into the optimal portfolio, whereas the rest of the futures contracts are less common with the October and November contracts used in April and May, respectively, and the January and April contracts in August and September. The point remains that a better model of seasonal volatility alters the optimal hedging strategy.

## CONCLUSION

The POTS model of Smith (2005) reveals that NYMEX natural gas futures prices exhibit time-to-maturity effects and strong seasonal variations in their price volatility. Their volatility rapidly increases in the last three months of trading and is higher for winter contracts than for spring and summer contracts. In addition, the POTS model reveals that the persistence of price shocks and, hence, the correlation between daily price changes in concurrently traded contracts exhibit substantial seasonal and cross-sectional variation. Specifically, price volatility is relatively high in two trading periods: early November to mid-January

for the January to April contracts and early May to September for all contracts maturing before the following March. Such volatility dynamics are closely related to the seasonal pattern of the U.S. natural gas storage in a way consistent with the theory of storage.

The resulting portrait of natural gas price volatility dynamics implies that a hedger should include into its portfolio a futures contract of at least three months to maturity to avoid high contract-specific volatility in nearby delivery months. In addition, the hedger should include the December contract to hedge against spot price risk during spring and summer months and either of the June, July, and August contracts while hedging during winter months. The optimal hedge ratio is high, ranging from 1.2 to slightly above 2.0, because the prices of the more distant futures contracts exhibit much smaller movement than the nearby contracts, even while they share much information regarding underlying market conditions.

These results suggest that the previous studies of the spot–futures price relationships and of the dynamics of natural gas futures prices are subject to misspecification bias in the variance structure of the disturbance terms in those studies' regression models. For example, models of commodity price dynamics should allow more flexible specifications, in particular, seasonal and cross-sectional variations in the factor loadings and the variance of the idiosyncratic error. Similarly, the analysis of spot–futures price relationships should allow seasonal and cross-sectional variation in the variance of the disturbance term. The assumption of a constant, more specifically, zero, basis is clearly inappropriate, for the cost of carry is not constant for a storable commodity with strong seasonality. The optimal hedging strategies implied by these misspecified models are noticeably ineffective with the variance of the portfolio return 18–80% higher than the minimum variance attained by the hedging strategy suggested by the estimated POTS model.

## BIBLIOGRAPHY

- Hamilton, J. D. (1994). State-space models. In R. F. Engle, D. L. McFadden (Eds.), *Handbook of econometrics* (Vol. 4). Amsterdam: Elsevier.
- Lien, D., & Root, T. H. (1999). Convergence to the long-run equilibrium: The case of natural gas markets. *Energy Economics*, 21, 95–110.
- Lucia, J. J., & Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5, 5–50.
- Manoliu, M., & Tompaidis, S. (2002). Energy futures prices: Term structure models with Kalman filter estimation. *Applied Mathematical Finance*, 9, 21–43.
- Modjtahedi, B., & Movassagh, N. (2005). Natural-gas futures: Bias, predictive performance, and the theory of storage. *Energy Economics*, 27, 617–637.
- Movassagh, N., & Modjtahedi, B. (2005). Bias in backwardation in natural gas futures prices. *Journal of Futures Markets*, 25, 281–308.

- Root, T. H., & Lien, D. (2003). Can modeling the natural gas futures market as a threshold cointegrated system improve hedging and forecasting performance? *International Review of Financial Analysis*, 12, 117–133.
- Samuelson, P. A. (1965). Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6, 41–49.
- Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *Journal of Finance*, 52, 923–973.
- Smith, A. (2005). Partially overlapping time-series model: A new model for volatility dynamics in commodity futures. *Journal of Applied Econometrics*, 20, 405–422.
- Suenaga, H. (2005). Analysis of spot-forward price relationships in the restructured electricity markets. Ph.D. Dissertation, University of California, Davis, California.
- Walls, W. D. (1995). An econometric analysis of the market for natural gas futures. *Energy Journal*, 16, 71–83.
- Williams, J. C., & Wright, B. D. (1991). *Storage and commodity markets*. New York: Cambridge University Press.