

Area of Research Excellence in Oil & Gas Management

Working Paper Series
2007-05
April

“Volatility Dynamics of NYMEX Natural Gas Futures Prices”
By Hiroaki Suenaga, Aaron Smith and Jeffrey Williams

ISSN 1834-6766
Area of Research Excellence in Oil & Gas Management
School of Economics and Finance
Curtin Business School
Curtin University of Technology
GPO Box U1987, Perth WA 6845 AUSTRALIA
Email: OGM@cbs.curtin.edu.au
Web: <http://www.cbs.curtin.edu.au/ogm>

VOLATILITY DYNAMICS OF NYMEX NATURAL GAS FUTURES PRICES

By

Hiroaki Suenaga¹, Aaron Smith² and Jeffrey Williams³

ABSTRACT

We examine the volatility dynamics of the NYMEX natural gas futures prices via the partially overlapping time-series (POTS) model of Smith (2005, *Journal of Applied Econometrics*). We show that, aside from a time-to-maturity effect, volatility exhibits two important features that are closely related to the seasonal cycle of US natural gas demand and storage. First, volatility is greater in the winter than in the summer. Second, 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.

¹ Hiroaki Suenaga is a Research Fellow, School of Economics and Finance, Curtin Business School, Curtin University of Technology and is also the corresponding author, e-mail: hiroaki.suenaga@cbs.curtin.edu.au; phone: +61-8-9266-4480.

² Aaron Smith is Assistant Professor, Department of Agricultural and Resource Economics, University of California, Davis and Member, Giannini Foundation.

³ Jeffrey Williams is the Daniel Barton DeLoach Professor, Department of Agricultural and Resource Economics, University of California, Davis and Member, Giannini Foundation.

1. INTRODUCTION

Natural gas demand in the US peaks between December and March, due to 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 and 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 illustrated¹ in Figure 1, which displays all 72 NYMEX futures prices on 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 the high marginal cost of natural gas production and the inelastic winter demand mean that shocks of even a small magnitude can cause a large price swing. At the same time, the high natural gas

¹ All figures and tables can be found in the Appendix, starting on page 27.

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.

In this paper, we examine the volatility dynamics of the 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, while 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 mis-specifying 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 illustrate 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 minimising the variance of portfolio return. We contrast this result with the results from previous studies, which ignore seasonal and cross-sectional variations in volatility of natural gas prices and which recommend the use of the adjacent delivery contract for hedging.

Much of the previous research on natural gas, not to mention other energy futures, has not accounted for 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) have applied one- and two-factor models to NYMEX futures prices and report that the estimated models exhibit strong seasonal variation while 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 subsequently realised 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 have reported that the NYMEX natural gas futures prices are downward biased in forecasting the subsequently realised spot price and interpreted this bias as representing a risk premium (Walls, 1995; Modjahedi and Movassagh, 2005; Movassagh and Modjahedi, 2005). 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 zero basis 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 existing futures markets. Williams and Wright (1991) show that the equilibrium futures prices derived from a dynamic rational expectation 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 variations in the

basis. In contrast, the POTS model avoids such misspecification by differencing out the price-level dynamics.

Aside from these modeling issues, most previous studies have utilised only a subset of price data available from organised 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. This practice is particularly common to tests of bias in futures prices because most energy futures contracts are defined in monthly blocks. 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 due to switching from one delivery month to another.

2. PARTIALLY OVERLAPPING TIME-SERIES MODEL

The partially overlapping time-series (POTS) model of Smith (2005) is a latent factor model of daily futures price changes. In this model, daily price change of a particular contract 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,

$$(1) \quad \Delta \mathbf{F}_t = \boldsymbol{\theta}_t \varepsilon_t + \boldsymbol{\lambda}_t \mathbf{u}_t$$

where $\Delta \mathbf{F}_t$ 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 ε_t is a latent factor with $E[\varepsilon_t] = 0$ and $E[\varepsilon_t^2] = 1$ and $\boldsymbol{\theta}_t$ is $n_t \times 1$ 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 of contract. That is,

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

For identification, we specify $E[\varepsilon_t u_{m,t}] = 0$ for all m and t , and $E[\Delta \mathbf{F}_t | \mathfrak{F}^{t-1}] = 0$ where \mathfrak{F}^{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.²

Following Smith (2005), we specify the conditional variance of ε_t by the GARCH(1,1) process,

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

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

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

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) through (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 price levels and not price changes. By modeling daily 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 bias could often be large in commodity price models, especially for commodities that exhibit such a strong seasonal pattern as natural gas (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 nonlinear volatility dynamics of natural gas futures prices resulting from peculiarities of the commodity. In contrast, conventional factor models (e.g., Schwartz (1997)) 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

² We also estimate 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.

the idiosyncratic errors with at most cross-sectional variation across contract delivery months.³

3. DATA AND ESTIMATION RESULTS

3.1 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 (BTu) of natural gas at the Henry Hub located in Louisiana. The market initially traded contracts as far as 1 year before the first calendar day of the delivery month, but has gradually extended this horizon to 6 years. Contracts were initially traded until seven business days prior to the first calendar day of the delivery month, but trading was later extended to 3 business days prior to the delivery month.

We analyse the daily change in the logarithm of settlement prices, and our sample spans January 2, 1991 to December 31, 2003. Because distant delivery contracts are not always actively traded, we drop contracts of more than 12 months to maturity. Excluding these observations, we have a sample of 40,618 prices among 175 contracts.

³ Suenaga (2005) illustrates 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.

3.2 Estimation Results

In applying the POTS model to the NYMEX natural gas futures price data, we specify the factor loadings and the variance of idiosyncratic error in (2) by the following trigonometric functions,

$$(2') \quad \begin{aligned} \theta_{m,t} = \theta^m(d) &= a_0^m + a_1^m d + \sum_{j=1}^5 \left(a_{2j}^m \sin\left(\frac{2\pi j d}{d_{\max}}\right) + a_{2j+1}^m \cos\left(\frac{2\pi j d}{d_{\max}}\right) \right) \\ \lambda_{m,t} = \lambda^m(d) &= b_0^m + b_1^m d + \sum_{j=1}^5 \left(b_{2j}^m \sin\left(\frac{2\pi j d}{d_{\max}}\right) + b_{2j+1}^m \cos\left(\frac{2\pi j d}{d_{\max}}\right) \right) \end{aligned}$$

where $d_{\max} = 365$. To capture seasonality we estimate one such function for each contract month i.e., for all $m = 1, \dots, 12$. Specified as in (2'), the two functions become more flexible as the number of trigonometric terms 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. By using only five trigonometric terms, we allow sufficient flexibility to capture seasonality and time-to-delivery effects, but we avoid excess sensitivity to outliers. We estimate the model by the method of Maximum Likelihood with the initial value obtained by the iterative approximate-EM method of Smith (2005).⁴

Table 1 summarises the coefficient estimates of the GARCH parameters. In the table, the coefficient estimates of α and β are 0.094 and 0.897. The sum of the two

⁴ 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 Kalman Filter, (2) maximise 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 the step 2.

coefficients, 0.991, is less than unity, indicating that the conditional volatility of the common underlying factor is highly persistent yet is stationary.

Figure 2 plots the unconditional variance of daily price changes of each of 12 contracts, which we compute 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, $\lambda_{m,t}$. The figure exhibits at least four interesting 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 effect (1965), indicates that shocks to spot prices are expected to dissipate somewhat over time. Second, 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 features 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 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 realised. Hence, price volatility increases during the May to late September trading period, but only for the contracts deliverable to 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 fair 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 3 illustrates this storage pattern. In the figure, the US 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 2 summarises the proportion of the variance explained by the common factor, ε_t . It indicates that the model explains, on average, about 81% of daily price variation. In general, spring and summer contracts are more closely related to the common factor than winter contracts. Figure 4 illustrates, for 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 percent 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 4 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 afterwards. 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 Figure 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 Figure 2), but a high proportion of their variation is captured by the common factor (see Figure 4). 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 price exhibits both the Samuelson effect and strong seasonality. 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 are unable to discern because their econometric models presumed a volatility specification that is too simple.

4 OPTIMAL HEDGING STRATEGY

In this section, we extend our analysis of the volatility dynamics of natural gas futures prices to investigate the implications for hedging. We consider a simple hedging strategy, 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 τ .

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

$$(4) \quad W_{t+k} = (S_{t+k} - S_t)Q - (F_{t+k,\tau} - F_{t,\tau})X = (\Delta S_{t+k} - \eta \Delta F_{t+k,\tau})Q$$

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

$$(5) \quad V[W_{t+k}] = (V[\Delta S_{t+k}] + \eta^2 V[\Delta F_{t+k,\tau}] - 2\eta \text{cov}[\Delta S_{t+k}, \Delta F_{t+k,\tau}])Q^2$$

which is minimised by,

$$(6) \quad \eta_{t,\tau}^* = \text{cov}[\Delta S_{t+k}, \Delta F_{t+k,\tau}] / V[\Delta F_{t+k,\tau}]$$

Substituting (6) into (5) yields the minimised variance,

$$(7) \quad V[W_{t+k} | \eta_{t,\tau}^*] = V[\Delta S_{t+k}] (1 - \rho_{t+k,\tau}^2) Q^2$$

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

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

$$(8) \quad \eta_{t,\tau}^* = \text{cov}[\Delta \ln S_{t+k}, \Delta \ln F_{t+k,\tau}] / \text{V}[\Delta \ln F_{t+k,\tau}]$$

$$(9) \quad \text{V}[r_{t+k} | \eta_{t,\tau}^*] = \text{V}[\Delta \ln S_{t+k}] (1 - \rho_{t+k,\tau}^2) Q^2$$

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

Three remarks should be made about these hedging strategies. First, many organised exchanges concurrently trade multiple contracts with different maturity dates. Thus, although the futures contract to be included into the portfolio, τ , is exogenous to the above strategy, it should be a decision variable for the hedger. The expressions (6) and (8) indicate that given the time of entry, t , and hedging horizon, k , the hedger should include into 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.

Second, the time of entry, t , and hedging horizon, k , should be also 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 practice would yield, at least on average, a negative return. The hedge ratios in (6)

and (8) 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, the expressions (6) through (9) 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 idiosyncratic error, yields an optimal hedge ratio that varies by τ , t , and k . In contrast, conventional models of commodity price dynamics determine factor loadings by the time to maturity of the contract and a small number of parameters defining the stochastic processes of the underlying factors. Due to this restrictive specification, the optimal hedge ratio implied by these models does not vary by contract delivery date. In particular, a simple one-factor mean-reversion model considered by Schwartz (1997) and the two-factor model by Manoliu and Tompaids (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 due to a simple variance structure assumed for the disturbance term.

4.1 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.⁵ In doing so, we use the nearby futures price as a proxy for the spot price. Based on our discussion in Section 3.2, 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 where t and k are predetermined and the hedger's decision is to choose the futures contract to be included into the portfolio, τ , and the hedge ratio. We find the optimal solution for the 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. We consider these two cases to determine whether very short holding periods are less sensitive to the modeling of price dynamics.

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

$$(10) \quad E[\Delta F_t \Delta F_t'] = \theta_t \theta_t' + \lambda_t \lambda_t'$$

With the assumption $u_{d,t}$ iid $\sim N(0, 1)$, the diagonal and off-diagonal elements are,

$$(11) \quad E[\Delta F_{m,t}^2] = \theta(d, m)^2 + \lambda(d, m)^2$$

⁵ 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 evaluations based on the *unconditional* variance, because our objective here is to draw a general implication about the need to model seasonality.

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

where $d = m - t$ and $d' = m' - t$. Using the expressions (11), the correlation between daily price changes of two futures contracts for delivery m_1 and m_2 is,

$$(12) \quad \rho_{m_1, m_2, t} = \frac{\theta(d_1, m_1) \theta(d_2, m_2)}{(\theta(d_1, m_1)^2 + \lambda(d_1, m_1)^2)^{\frac{1}{2}} (\theta(d_2, m_2)^2 + \lambda(d_2, m_2)^2)^{\frac{1}{2}}} = \pi_{m_1, t} \pi_{m_2, t}$$

where $d_i = m_i - t$ and $\pi_{m_i, t} = \theta(d_i, m_i) (\theta(d_i, m_i)^2 + \lambda(d_i, m_i)^2)^{-\frac{1}{2}}$ is the square root of the share of the total variance of contract i explained by the common factor. Equation (12) indicates that the variance of the portfolio return is minimised 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 minimised variance are,

$$(13) \quad \eta_{m_1, m_2, t}^* = \frac{\theta(d_1, m_1) \theta(d_2, m_2)}{\theta(d_2, m_2)^2 + \lambda(d_2, m_2)^2}$$

$$(14) \quad V[W_{t+1} | \eta_{m_1, m_2, t}^*] = (\theta(d_1, m_1)^2 + \lambda(d_1, m_1)^2) (1 - \rho_{m_1, m_2, t}^2) Q^2$$

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,

$$(15) \quad E[\Delta F_{t_1, t_2} \Delta F_{t_1, t_2}'] = \sum_{s=t_1}^{t_2} (\theta_s \theta_s' + \lambda_s \lambda_s')$$

with its elements,

$$(16) \quad E[\Delta F_{m, t_1, t_2}^2] = \sum_{s=t_1}^{t_2} \theta(m-s, m)^2 + \lambda(m-s, m)^2$$

$$E[\Delta F_{m, t_1, t_2} \Delta F_{m', t_1, t_2}] = \sum_{s=t_1}^{t_2} \theta(m-s, m) \theta(m'-s, m') \quad \text{for } m, m' > t_2, m \neq m'.$$

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

$$(17) \quad \rho_{m_1, m_2, t_1, t_2} = \frac{\sum_{s=t_1}^{t_2} \theta(m_1-s, m_1) \theta(m_2-s, m_2)}{\left(\sum_{s=t_1}^{t_2} \theta(m_1-s, m_1)^2 + \lambda(m_1-s, m_1)^2 \right)^{\frac{1}{2}} \left(\sum_{s=t_1}^{t_2} \theta(m_2-s, m_2)^2 + \lambda(m_2-s, m_2)^2 \right)^{\frac{1}{2}}}$$

The optimal hedge ratio and the associated minimum variance are,

$$(18) \quad \eta_{m_1, m_2, t_1, t_2}^* = \frac{\sum_{s=t_1}^{t_2} \theta(m_1-s, m_1) \theta(m_2-s, m_2)}{\sum_{s=t_1}^{t_2} \theta(m_2-s, m_2)^2 + \lambda(m_2-s, m_2)^2}$$

$$(19) \quad V[W_{m_1, m_2, t_1, t_2} | \eta_{m_1, m_2, t_1, t_2}^*] = \sum_{s=t_1}^{t_2} (\theta(m_1-s, m_1)^2 + \lambda(m_1-s, m_1)^2) (1 - \rho_{m_1, m_2, t_1, t_2}^2) Q^2$$

Equations (13) and (18) 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.

4.2 Optimal Hedging Strategy for Natural Gas

Figures 5 through 7 illustrate the futures contract included into the optimal portfolio, the optimal hedging ratio, and the minimum variance attained by the optimal portfolio for each of the two hedging horizons. First, for a daily hedging strategy, 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 contracts are often 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 into 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 minimising contract switches gradually from the August to winter and from the December to summer contracts, respectively.

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 year, the optimal portfolio includes the July contract and then it switches to the August contract. Contracts for earlier maturity are not used due to their high idiosyncratic volatility. This high idiosyncratic volatility

represents the Samuelson effect for the March contract whereas it is due to 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 the maturity effect. This transition is rather quick with the August contract replaced by the December contracts 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 contract are optimal only for a short period because price movements of these two contracts are contract-specific due to the low inventory applicable during those two months.

Figure 6a plots the optimal hedge ratio as a function of the date within the gas year. As shown in (13), 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 6b 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 Tompaides (2002).⁶ Two observations become clear with this comparison of the two hedging strategies. First, the optimal hedge ratio is substantially lower for the portfolio including the second position 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 as the contract approaches maturity. This is, again, because the share of the price variance explained by the common factor decreases due to the Samuelson effect. These results are peculiar to the mean-reversion process assumed for the underlying factors in the models of Manoliu and Tompaides (2002). This particular stochastic process assures that the factor loading increases monotonically as the contract approaches its maturity date whereas the variance of the idiosyncratic error by specification is 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.

⁶ See Tables 4 and 5 on pages 38 and 39 of Manoliu and Tompaides (2002). Numbers in Figure 6b were calculated by using the volatility estimates of the POTS model in (13).

Figure 7a 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. 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 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.

Figure 7b compares the optimal hedging strategy implied by the POTS model with the straightforward strategy of utilising the second-position contract. The figure indicates that the strategy based on the POTS model attains the portfolio variance that is about 15 to 45 percent below the variance of the portfolio utilising 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 optimal hedging strategies for a monthly horizon yield essentially the same results as those for a daily horizon. 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

contract used in April and May, respectively, and the January and April contracts in August and September. The optimal hedging ratio and the minimum variance attained by the optimal contract are also similar to those for daily hedging strategy. They are almost identical to the monthly averages of their corresponding values for the 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, due to a small variance of the futures contract included in the portfolio.

5. CONCLUSION

We examine the volatility dynamics of NYMEX natural gas futures prices using the partially overlapping time-series (POTS) model of Smith (2005). The estimated POTS model reveals that the NYMEX natural gas futures prices exhibit time-to-maturity effects and strong seasonal variations in their price volatility – volatility rapidly increases in the last three months of the trading period and is higher for winter contracts than for spring and summer contracts. In addition, our analysis 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 US natural gas storage in a way consistent with the theory of storage.

The depicted portrait of natural gas price volatility dynamics implies that a trader in need of hedging price risk should cross hedge with a futures contract of at least three

months to maturity to avoid high contract-specific volatility in nearby contracts. In addition, they should include in their portfolio the December contract to hedge against spot price risk during spring and summer months and either of the June, July, and August contract in winter months. The optimal hedge ratio is high, ranging from 1.2 to slightly above 2.0, because the price of the futures contract exhibits much smaller movement than the nearby contract while they share much information regarding underlying market conditions.

These results suggest that the previous studies of the spot-futures price relationships and the dynamics of natural gas futures prices are subject to misspecification bias in the variance structure of the disturbance terms in their regression models. For example, models of commodity dynamics should allow more flexible specifications, in particular, seasonal and cross-sectional variations in the factor loadings and the variance of the idiosyncratic error. Also, 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 in demand and/or supply. The optimal hedging strategies implied by these misspecified models are noticeably ineffective with the variance of the portfolio return 18 to 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 Engle, R. F., and McFadden, D. L. eds. *Handbook of Econometrics*, Vol. 4. Elsevier, Amsterdam.
- Lien, D., and Root, T. H. (1999). "Convergence to the long-run equilibrium: The case of natural gas markets," *Energy Economics*, 21: 95-110.
- Lucia, J. J., and Schwartz, E. S. (2002). "Electricity prices and power derivatives: Evidence from the Nordic Power Exchange," *Review of Derivatives Research*, 5: 5-50.
- Manoliu, M., and Tompaidis, S. (2002). "Energy futures prices: Term structure models with Kalman filter estimation," *Applied Mathematical Finance*, 9: 21-43.
- Movassagh, N., and Modjtahedi, B. (2005). "Bias in backwardation in natural gas futures prices," *Journal of Futures Markets*, 25: 281-308.
- Modjtahedi, B., and Movassagh, N. (2005). "Natural gas-futures: Bias, predictive performances, and the theory of storage," *Energy Economics*, 27: 617-637.
- Root, T. H., and 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: 117-131.
- 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., and Wright, B. D. (1991). *Storage and Commodity Markets*. Cambridge University Press, New York.

APPENDIX

Table 1: Maximum Likelihood Estimates of GARCH Parameters

	Coefficient	Std. Err.	<i>t</i> -ratio
GARCH Parameters			
α	0.0936	0.0321	2.9178
β	0.8972	0.0214	41.8655
Log Likelihood			185844.9
BIC			4.4997

Table 2: Proportion of the Variance Explained by a Common Factor

Overall	
	0.8131

By contract	
1	0.7613
2	0.7750
3	0.7523
4	0.8319
5	0.8469
6	0.8584
7	0.8622
8	0.8756
9	0.8612
10	0.8380
11	0.8140
12	0.7465

Figure 1: NYMEX natural gas settlement prices as of April 1, 2003

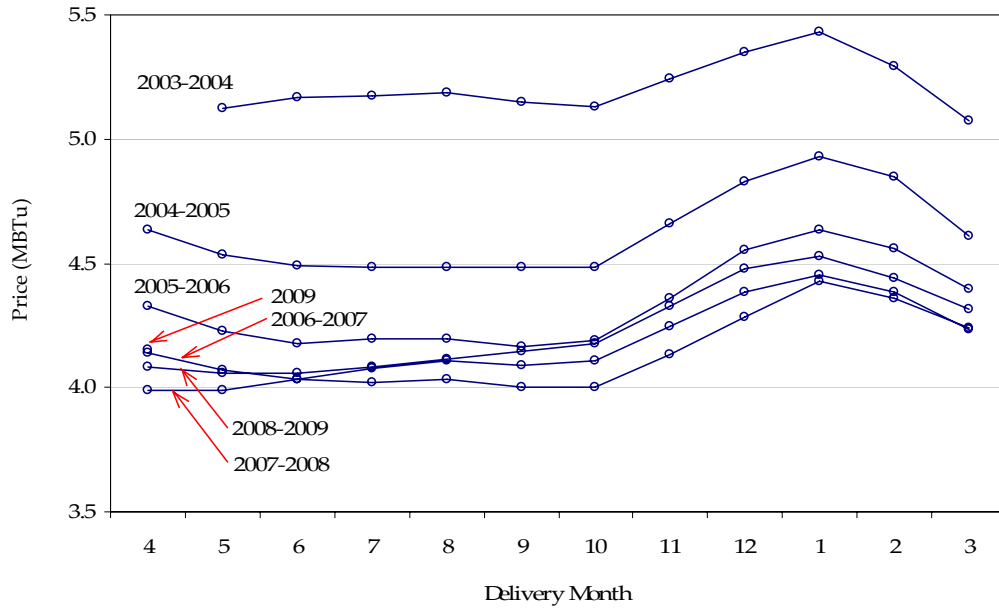


Figure 2: Variance of daily log price changes implied by the estimated POTS model




Figure 3: US natural gas underground storage – Monthly average of working gas for 1976-2005




Figure 4: Share of price variation explained by the common factor

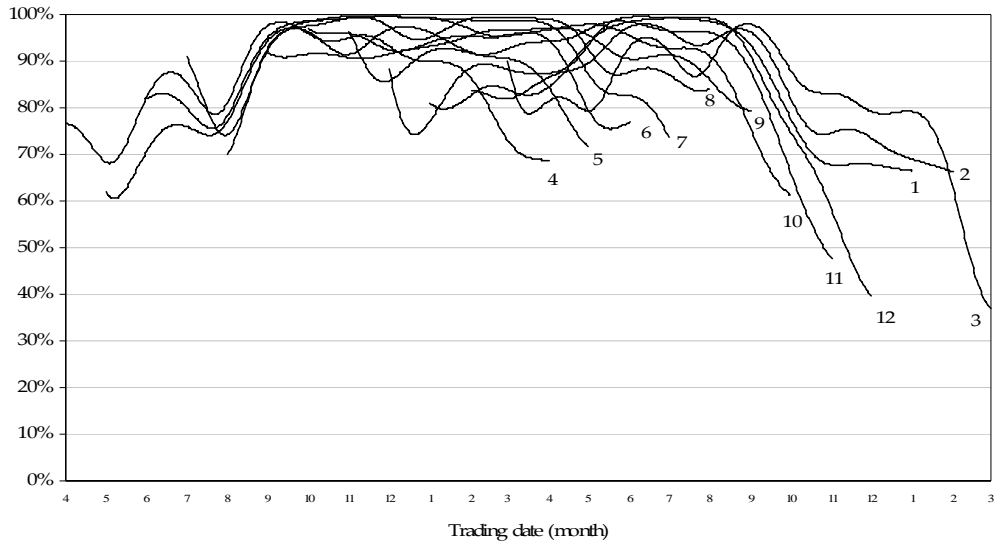


Figure 5: Delivery month of the futures contract included in the optimal portfolio


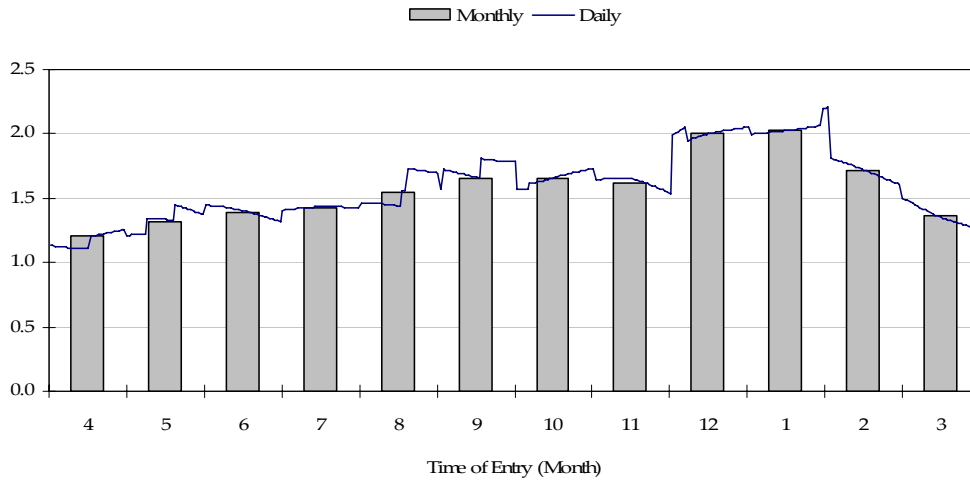


Figure 6: Optimal hedge ratio

(a) Portfolio suggested by the estimated POTS model



(b) Portfolio including the second position contract





Figure 7: Minimum variance attained by the optimal portfolio

(a) Relative to the variance of unhedged portfolio



(b) Relative to the variance of portfolio including the second position contract

