FORECASTING PERFORMANCE OF FUNDAMENTAL NATURAL GAS PRICE MODELS, HEDGING STRATEGIES, AND THE AVERAGE COST OF GAS: A STUDY OF THE U.S. NATURAL GAS MARKET

Scott C. Linn and Zhen Zhu*

We propose and estimate fundamental models for natural gas prices. We compare how well these models, as well as univariate statistical time series models of NG prices and the NYMEX futures price for natural gas, forecast spot gas prices. We find that a univariate time series model that incorporates fundamental variables related to production, storage, weather, and aggregate output performs best when NG prices are falling. In contrast, when prices are rising, a VAR specification with multiple fundamental endogenous and exogenous variables gives the best predictions for time horizons of either 6, 9, or 12 months, while the futures price gives the best predictions for a 3-month horizon. We also examine the average gas cost to a user who implemented either an Always Hedge, Never Hedge, or a Mixed Strategy based upon the price forecast. The Mixed Strategy utilizes the spot price forecasts based upon the alternative forecasting models. We find that during the falling price phase, but irrespective of the forecast model utilized, the Mixed Strategy always produces an average cost that is less than or equal to the strategy of Always Hedging. The same is true during a rising price phase for the 9- and 12-month time horizons, but during the 3- and 6- month time horizons the policy of always hedging dominates. However, we also find that during a falling price phase the absolute least cost strategy is to never hedge and that this is also the dominating strategy during the rising price phase for 9- and 12-month out horizons.

Scott C. Linn (the corresponding author) is a professor in the Division of Finance in the Michael F. Price College of Business at the University of Oklahoma. Contact information: University of Oklahoma, Michael F. Price College of Business, Division of Finance, 205A Adams Hall, University of Oklahoma, Norman, OK 73019. E-mail: slinn@ou.edu.

Zhen Zhu works for C.H. Guernsey & Company and is also an adjunct professor in the Department of Economics, College of Business, at the University of Central Oklahoma. Contact information: Department of Economics, College of Business, University of Central Oklahoma, Edmond, OK 73034. E-mail: Zhen.Zhu@chguernsey.com.

Acknowledgement: We are grateful to the Financial Research Institute of the University of Missouri for financial support of this research and to an anonymous referee for many helpful comments. We also thank C. Fernando, M. Turac, and seminar participants at the University of Oklahoma for comments on an earlier draft of the paper. Scott Linn received support from the John and Mary Nichols Faculty Fellow Program and the Harold S. Cooksey Endowment and thanks the sponsors. Zhen Zhu also gratefully thanks the Joe Jackson College of Graduate Studies and Research at the University of Central Oklahoma for financial support.

Keywords: energy price modeling, risk management, statistical models

JEL Classification: Q40, Q41

Most observers agree that commercial users of natural gas, including electricity producers which use natural gas as an input, often hedge their natural gas requirements (Fitzgerald and Pokalsky 1995; Clewlow and Strickland 2000; Eydeland and Wolyniec 2003). Many also hold the view that natural gas forward or futures markets are inefficient in the sense that arbitrage opportunities not present in futures markets for financial instruments may be present in natural gas forward or futures markets due possibly to informational inefficiencies (Leong 1997; Murry and Zhu, 2004). Finally, some who are close to this market have suggested that many ostensible natural gas hedgers actually incorporate an element of speculation in their trades.

If inefficiencies exist, the natural gas user might gain from selectively choosing when and when not to hedge. The ability of the natural gas user to implement a mixed strategy in which the user sometimes hedges and sometimes selects to trade on the spot market is predicated on the availability of an accurate model of the natural gas spot price that can be used to forecast future spot prices. The purpose of this study is to examine whether strategies based upon predictions from models of the level of the spot natural gas price can be used to exploit inefficiencies in the natural gas (NG) market, thereby producing average user costs below those implied by a strategy of never hedging or a strategy of always hedging using the NYMEX Natural Gas futures contract.

We propose and estimate fundamental models for natural gas prices using as a basis the theory of price determination for storable commodities. There are two basic approaches used in modeling the behavior of commodity prices. The first, and the one we follow in this study, articulates the model in terms of proxies for the fundamental determinants of supply and demand. The second, which is generally less concerned about the level of prices but more concerned about the volatility of prices, articulates price behavior in terms of a univariate generalized stochastic process, in which the structural form of the model attempts to account for the fundamental determinants without directly specifying their influence or behavior. The latter models generally serve as a foundation for the valuation of derivative instruments in which volatility is of primary concern and consequently with short-range behavior.

We then compare how well these models, as well as univariate statistical time series models of NG prices, forecast spot gas prices. The comparison includes an analysis of the predictive performance of each model including an analysis of the predictive performance of the NG futures price for the NYMEX traded contract. We examine two distinct periods, one during which NG prices were falling and another during which prices were rising. We find that a univariate time series model that incorporates fundamental variables related to production, storage, weather, and aggregate output performs best in a root mean square error sense among all the models examined when NG prices are falling. Likewise, when prices are rising, a VAR specification with multiple fundamental endogenous and exogenous variables gives the best predictions for time horizons of either 6, 9 or 12 months, while the futures price gives the best predictions for a 3-month horizon.

We also examine the average gas cost to a user who implemented either an Always Hedge, Never Hedge, or a Mixed Strategy based upon strategically selecting to hedge or not to hedge based upon the price forecast. The Mixed Strategy utilizes the spot price forecasts based upon the alternative forecasting models. Several different time horizons are assumed. We find that during a falling price phase, but irrespective of the forecast model utilized, the Mixed Strategy always produces an average cost that is less than or equal to the strategy of always hedging. The same is true during a rising price phase for the 9- and 12-month time horizons, but during the 3 and 6 month time horizons the policy of always hedging dominates. However, we also find that during a falling price phase the absolute least cost strategy is to never hedge and that this strategy dominates during a rising price phase for 9- and 12-month out horizons.

Our results shed light on the potential effectiveness of fundamental models of the NG spot price for natural gas users intent on minimizing their average cost of gas. The results suggest that, if the price phase cycle is not clear, a mixed strategy can be an effective tool for minimizing cost.

Natural gas (NG) prices in the North American natural gas market have reached unheard of levels in recent years, not to mention the increases in price volatility that have also been observed. The commercial demand for natural gas is not, however, expected to abate.[1] Accurate predictions of the future spot price can play an important role in cost minimization strategies of natural gas users who might otherwise select to always hedge using the futures market. Thus, accurate predictions potentially help to forestall the exacerbation of financial fragility among commercial users of natural gas brought on by unexpectedly high NG prices.[2] Understanding the nature and determinants of natural gas prices is therefore of important practical interest to both commercial users as well as policy-makers. However, academic studies of natural gas price behavior are limited. In addition, these studies have tended not to focus on articulating and modeling the fundamental determinants of spot natural gas price movements. Instead, they have focused on generalized statistical processes, more often than not where concern is with the pricing of various types of derivative instruments where a measure of volatility is the key ingredient. (An excellent example of the latter approach is Schwartz 1997.) In contrast our focus is on mean cash flow effects. We believe both approaches have merit and provide complementary guidance particularly in light of the growing evidence that many firms are believed to engage in selective hedging, essentially taking positions based upon their expectations about the future level of the price (Stulz 1996; Brown, Crabb, and Haushalter 2005; Knill, Minnick, and Nejadmalayeri 2005). Our study

1.Electricity generating utilities are major users of natural gas as are chemical and fertilizer producers. The Energy Information Administration for instance projects that natural gas fired electricity generating capacity will increase dramatically over the next 20 years accounting for 31% of natural gas demand by 2025 as compared with 23% in 2003 (EIA Annual Energy Outlook 2005: http://www.eia.doe.gov/oiaf/aeo/index.html. The EIA expects that most new electricity generating capacity will be fueled by natural gas.

2. One has only to consider the California electricity crisis of 2000-2001 to appreciate the manifold problems that can arise.

contributes to the literature focusing on the behavior of natural gas prices, and in particular on the forecastability of spot natural gas prices in the U.S. natural gas market. Equally important is the question of the reliability of the NG futures price as a predictor of the spot price.

We begin with a brief review of the theory of spot price determination within the context of a model that allows storage. The model highlights the fundamental determinants of storable commodity prices. We then turn to a discussion of the natural gas price series analyzed in the paper and the fundamental variables we employ as proxies for the determinants of prices. This is followed by an analysis of the time-series properties of spot natural gas prices and the fundamental variables. We then present estimation results for models of natural gas prices and also present an examination of price dynamics for a selection of multivariate models. We then turn to a comparison of the forecasting ability of both univariate as well as multivariate models of the natural gas price and the forecasting ability of the NG futures price. We conclude with an analysis of the out-of-sample cost of several hedging strategies in which a forecasting model for the price is used as the basis for the strategy. The strategies are predicated on the user seeking the minimum average cost of gas. We end with a summary of our findings and conclusions about the usefulness of natural gas price models and the predictive quality of the NG futures price.

I. A SIMPLE THEORY OF STORABLE COMMODITY PRICES

The theory of price determination for storable commodities has a long history, including early work by Gustafson (1958), Muth (1961), and Samuelson (1971). Recent contributions, from which we borrow heavily in the following discussion, are Deaton and Laroque (1992, 1996). (See also the work by Bresnahan and Suslow 1985; Williams and Wright 1991; and Chambers and Bailey 1996.)

Let the price of the commodity at time t be given by p_t. Assume there are two classes of agents in the model, producers-consumers and speculators. The excess demand of producers-consumers is determined by the current price. Speculators are inventory holders who engage in carrying the commodity forward until the next period.

We begin by assuming that inventory holding is not allowed. Define $Q = D(p_t)$ as the excess demand function for the commodity and let z_t represent the net random production (= net excess demand) of the commodity in period t. The equilibrium price will therefore be given by $D(p_t) = z_t$. We let $p_t = P(z_t)$ represent the inverse demand function for the commodity. Consider the simplest case in which z_t is independently and identically distributed over time. The price $p_t = P(z_t)$ will therefore also be iid as it will be determined by the random production in period t and the shape of the excess demand function. Prices will therefore be uncorrelated across time. Actual natural gas prices, however, exhibit autocorrelation. Natural gas of course is a storable commodity, so the model needs to be modified to account for inventories if it is to provide a reasonable description of the natural gas price formation process.

As we will see, the introduction of inventory holding can result in prices no longer being iid.

Assume inventory may be held but at a cost. Specifically, inventory deteriorates in value at the known rate $\delta > 0$ so that a unit of the commodity stored at t yields (1- ä) units at t+1. Storage of natural gas exhibits this feature. Natural gas in the United States is typically stored in either salt or aquifer caverns. Costs are incurred in injecting and withdrawing gas as well as for carrying the inventory (Sturm 1997).

For simplicity, inventory holders are assumed to be risk-neutral expected profit maximizers who can borrow and lend at the known rate of interest $r > 0$. Define $E(p_{t+1})$ as the expected t+1 price as perceived by speculative inventory holders at t. The inventory holder's present value of expected profit from the decision to place in inventory today I_t units of the commodity is given by

$$\begin{aligned} \Pi_t &= \pi_t I_t \\ &= \left\{ \frac{1}{1+r}(1-\delta)(E(p_{t+1})) - p_t \right\} I_t \end{aligned} \tag{1}$$

Value maximization implies that in equilibrium

$$\begin{aligned} I_t &= 0 \; if \; \pi_t < 0 \\ I_t &\geq 0 \; if \; \pi_t = 0. \end{aligned} \tag{2}$$

In other words, when the present value of expected profits is negative, no inventory is held. If there is a positive expected profit from holding the marginal unit, prices will be bid up to the point where the marginal profit on the next unit is equal to zero. The market clearing condition is given by

$$z_t + (1-\delta)I_{t-1} - I_t = D(p_t) \tag{3}$$

where the left hand side of (3) represents net supply and the right hand side represents demand where the quantity $(1 - \delta)I_{t-1}$ equals the inventory that survives after being held from t-1 to t. The above implies the price must satisfy the following condition

$$p_t = \max\left\{ \frac{1}{1+r}(1-\delta)(E(p_{t+1})), P(z_t + (1-\delta)I_{t-1}) \right\} \tag{4}$$

The first term on the right hand side applies if inventory is held and equals the expected payoff per unit of inventory held. The second term applies if no inventory is held from t to t+1, in which case all inventory carried into t and all net production z_t are consumed. Deaton and Laroque (1992) show that the rational expectations equilibrium price for this setting can be characterized as

$$p_t = f(I_{t-1}, z_t) \tag{5}$$

where the function $f(I_{t-1}, z_t)$ depends upon the fundamental parameters of the system, δ, r, the shape of the inverse demand function and the probability distribution of random net production (demand) z_t. For our purposes it is enough to recognize that inventory storage and production (excess demand) are the key ingredients.

One important result of the model is that prices can exhibit autocorrelation when storage is allowed. When prices are low, there is no storage, but low prices may in turn lead to demand for speculative storage, causing the price to be bid up. In turn, higher prices lead to liquidation of inventory stocks, causing the price to fall. Thus, inventories should be more common than uncommon if autocorrelation in prices is to be observed. In fact, the storage of natural gas in the United States is common. Aside from the mere presence of speculative inventory holders, another potential source of autocorrelation is the behavior of excess demand. The assumption that z_t is iid is potentially unrealistic. Deaton and Laroque (1996) speculate that z_t may exhibit autocorrelation and explore the implications of relaxing the iid assumption. They show that for small values of δ, not only does autocorrelation in z_t lead to autocorrelation in prices, but the autocorrelation in prices is greater than the autocorrelation in excess demand. Further, if the autocorrelation in excess demand is less than 1 (the stationary case), so also will be the autocorrelation in prices, for realistic values of r. Taken together, the presence of speculative inventory holders and/or autocorrelation in excess demand may induce autocorrelation in the price process.

The model illustrated above is intended to act as a framework for building a statistical model of the natural gas price process. Because our focus is on the forecasting ability of such statistical models, we do not explicitly test the economic model but instead use it as a guide for defining variables that should theoretically influence the price and that are observable by a decision maker operating in the natural gas spot and futures markets.

II. DETERMINANTS OF NATURAL GAS PRICES

A. Data Definitions and Preliminary Statistics

We examine natural gas spot prices for delivery at the Henry Hub. The spot price data are from the database maintained by Platts, publisher of the industry news and data source *Platts Gas Daily*.[3] The Sabine Pipeline Hub at Henry, Louisiana, is the designated delivery location for the NYMEX natural gas futures contract. We examine the time-series of weekly averages of daily mid-point spot prices.[4] The price series spans the period January 4, 1991, through June 7, 2002.

3. The Gas Daily Henry Hub spot price is generally regarded as the best spot gas prices available.
4. The theory developed in Deaton and Laroque (1996) is a theory of real prices. Our focus is on the prediction of nominal prices, not with a direct test of the theory. The wll-known problems with forecasting inflation lead us to focus exclusively on nominal gas prices rather than attempting to construct models of both real prices and the appropriate inflation rate.

Figure 1. Henry Hub Natural Gas Prices.

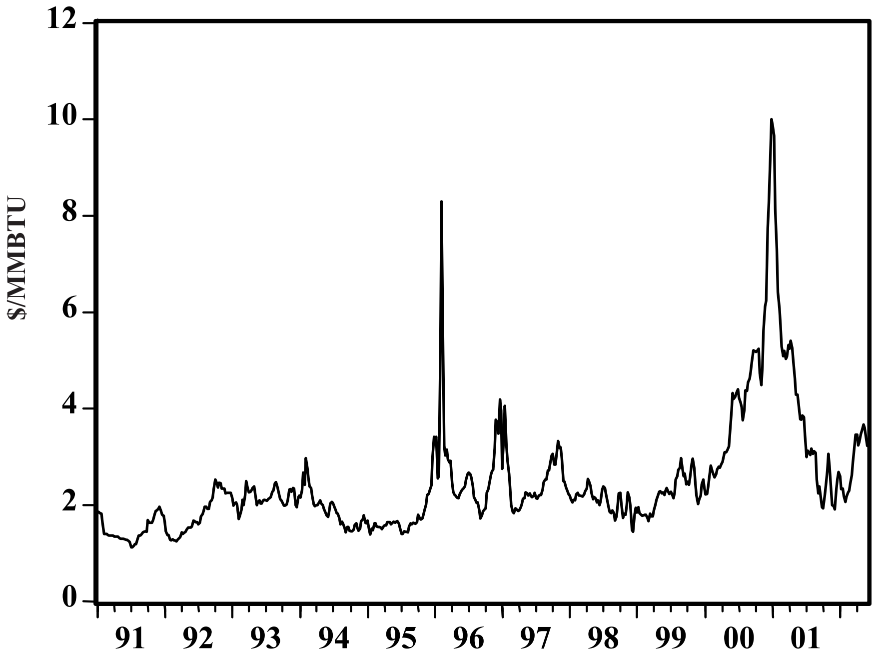

Weekly averages of daily mid-point spot prices for gas delivered at the Henry Hub (Sabine, LA).
Source: Platts

Weekly data are examined for two reasons. First, we develop several models of the fundamental determinants of natural gas prices for which the fundamental variables are available only on a weekly frequency. The second reason for modeling weekly data is that most users who are constructing forecasts for the purpose of hedging will be concerned with intermediate to long-term horizons. The usefulness of a daily price model for forecasting purposes is not likely to add much when the forecast horizon is 3 to 12 months and may in fact lead to larger forecast errors than a model that smoothes daily behavior. A plot of the natural gas price series is presented in Figure 1. The data exhibit the now well known episodes from the mid-1990s and early 2000s, during which prices surged as well as general seasonal cycles. In addition, the data suggest an upward trend in the price over the sample period.

The statistical models are based upon publicly available information. In addition we also examine the predictive ability of the NG futures price.

Table 1 presents descriptive statistics for the data utilized in the analyses. Prices are stated in terms of MMBTU.[5] Platts obtains within day transaction prices and the quantities traded. We utilize the daily mid-point price in constructing our

5. A cubic foot of natural gas on average gives off about 1,025 to 1031 Btu's. One Btu (British Thermal Unit) is the amount of heat required to raise the temperature of one pound of water from 60 to 61 degrees Fahrenheit at normal atmospheric pressure (14.7 pounds per square inch).

Table 1. Descriptive Statistics

Variable	Sample Period	Mean	Std. Dev.	Skewness	Kurtosis	J-B
Price	1/4/91-6/7/02	2.47	1.216	2.85	11.09	3868.8
				(0.00)	(0.00)	(0.00)
Quantity	1/29/99-6/7/02	4101	2114	0.22	-1.188	11.837
				(0.227)	(0.0016)	(0.003)
Oil Price	1/4/91-5/31/02	20.92	4.755	0.689	0.364	50.45
				(0.00)	(0.072)	(0.00)
Rig Count	1/4/91-6/7/02	515	185.7	1.079	0.72	128.93
				(0.00)	(0.0035)	(0.00)
Storage	1/7/94-6/7/02	-0.636	88.056	-0.80	-0.462	51.12
				(0.00)	(0.0495)	(0.000)
IP	1/4/91-6/7/02	1.005	0.074	-0.26	-1.37	53.60
				(0.01)	(0.00)	(0.00)

Note: p-values for tests of zero skewness, excess kurtosis and Jarque-Bera normality tests are in parentheses.
Price: Weekly average of natural gas spot price at Henry Hub (daily mid-point)
Quantity: Weekly average of daily volumes (in 000's MMBTU) of gas traded at Henry Hub
Oil Price: Weekly average of daily WTI crude oil price
Rig Count: Baker-Hughes weekly U.S. natural gas rig count
Storage: Weekly injection/drawdown of underground gas storage
IP: Gas weighted industrial output index, monthly data extrapolated into weekly
Data sources: Price and quantity data are from archives maintained by *Platts*, Rig Count is from Baker-Hughes, Storage data is from the AGA and the EIA, and Oil Price and IP are from the EIA.

weekly series.[6] The weekly data series are constructed by averaging the daily mid-point prices over the days of the week, by week. Henceforth, we will refer to these values as the weekly prices. The mean weekly price for the sample period is \$2.47 with a standard deviation of \$1.216. The price data exhibit positive skewness and excess kurtosis relative to normality. The Jarque-Bera test soundly rejects normality for the price data.[7]

6. The *Gas Daily* does not prepare a volume-weighted price.
7. The Jarque-Bera statistic is used to test the hypothesis that a given set of data is drawn from a normal distribution. The test statistic measures the difference of the skewness and kurtosis of the series with those from the normal distribution. The statistic is computed as:

$$JB = \frac{N-k}{6}\left(S^2 + \frac{1}{4}(K-3)^2\right)$$

where S is the skewness, K is the kurtosis, and k represents the number of estimated parameters used to create the series (Jarque and Bera, 1987). Under the null hypothesis of a normal distribution, the Jarque-Bera statistic is distributed as χ^2 with 2 degrees of freedom.

The fundamental predictive variables utilized in the analysis include the following, where the data sources are indicated in parentheses:

1. Quantity: Weekly average of daily volumes of gas traded at Henry Hub in 1000's MMBTU (*Platts*)
2. Oil Price: Weekly average of daily WTI crude oil price in $/barrel (*Energy Information Administration, U.S. Government*)
3. Rig Count: Baker-Hughes weekly U.S. natural gas rig count (*Baker-Hughes*)
4. Storage: Weekly injection/drawdown of underground gas storage (*American Gas Association; Energy Information Administration after 4/26/02*)
5. IP: Gas weighted industrial output index, monthly data extrapolated into weekly (*Energy Information Administration, U.S. Government*)

As already mentioned, the total of all quantities traded within a day is recorded by Platts. We construct the weekly series of quantities traded by averaging the daily quantities. Quantity is measured in terms of 1000's MMBTU. Quantity is used as a proxy for demand. The variable Oil Price is obtained from records maintained by the Energy Information Administration of the U.S. government. The Oil Price data series are the weekly averages of the daily West Texas Intermediate crude oil price. Oil Price serves as a control for the price of substitute energy products.[8] Baker-Hughes reports the number of rigs extracting natural gas for the United States on a weekly basis. We label this series Rig Count. During the sample period the American Gas Association produced a report issued weekly, indicating the amount of natural gas in storage as of the prior Friday. The report has been prepared by the Energy Information Administration of the U.S. government since April 26, 2002. The report shows the level as well as the net injection/drawdown from the total pool of gas. The variable Storage is the series of net injection/drawdown values obtained from the records of the AGA and the EIA. Storage serves as a proxy for the net change in inventory. The variable IP is also obtained from the Energy Information Administration. The EIA computes a gas-weighted index of industrial production for the United States on a monthly basis. The method used to compute the index is to weight industrial production across sectors based upon the relative consumption of natural gas by each sector. Therefore, the variable reflects general economic activity stated in terms of relative gas usage and serves as a general control for other factors potentially impacting natural gas supply and demand conditions. Price and Quantity are treated as endogenous variables. The remaining variables are treated as exogenous in some models and as endogenous in others. The Jarque-Bera statistics presented in Table 1 indicate that we reject the null hypothesis of normality for every variable.

8. Another substitute for natural gas is liquid natural gas (an imported product). However, the contribution of LNG to the overall natural gas market in the United States is relatively insignificant. According to the Energy Information Administration of the United States, LNG currently accounts for only about 1% of total U.S. consumption (see http://www.eia.doe.gov/oiaf/analysispaper/global/uslng.html). We therefore for parsimony, chose not to include the LNG price in the models we estimate.

B. Natural Gas Prices and the Weather

Demand for natural gas, and consequently the price, is determined in part by the seasonal weather cycle. The gas delivered at the Henry Hub must travel to the Hub via a specific set of natural gas pipelines. The weather in the areas serviced by the pipelines will naturally influence net demand and the price. Peak demand periods occur during the winter and the summer. Summer months are generally classified in the industry as June, July, and August and winter months as November to March.[9] The Shoulder months are those during which the weather is transitioning between the peak periods. The Shoulder Month 1 months are April and May, and Shoulder Month 2 months are September and October. The behavior of prices as seen in Figure 1 is consistent with these seasonal effects. The average price over the winter season equals 2.63 (standard deviation = 1.55) while the average price over the summer season equals 2.30 (standard deviation = .82).

In the models presented later, weather variables are included as controls for seasonal effects. Two weather related variables are introduced. The variable CDD Cooling Degree Days) is a daily temperature dependent numeric value which proxies for the propensity to use energy to "cool." As defined and used in the energy industry, CDD equals the daily average temperature minus 65 if the daily average temperature is higher than $65° F$. Likewise the variable HDD (Heating Degree Days) is a daily temperature dependent numeric value that proxies for the propensity to use energy to "heat." HDD equals 65 minus the daily average temperature if the daily average temperature is lower than $65° F$.[10] Daily temperature observations were obtained from regional Federal climate centers monitoring weather in the primary areas served by pipelines connecting at the Henry Hub. The cities included are Atlanta, Baton Rouge, Chicago, Dallas, Los Angeles, Little Rock New York, Phoenix, Philadelphia, and St. Louis. We prepared a weather index from these source data by computing the average temperature across the listed cities by day.[11] CDD and HDD were then computed from the weather index using the threshold temperatures indicated above.

C. Tests for Stationarity

The models we estimate are time-series formulations. We begin with an examination of the stationarity of each series using the Augmented Dickey-Fuller

9. The demand for gas during the summer has increased in recent years due to the industry bringing on-line more gas-fired electricity generating capacity. The Energy Information Administration projects that natural gas fired electricity generating capacity will continue to increase dramatically over the next 20 years (EIA Annual Energy Outlook 2005: http://www.eia.doe.gov/oiaf/aeo/index.html).

10. The definitions we employ for computing CDD and HDD are those commonly used by practitioners. Details on CDD and HDD along with time-series histories can be found on the website of the U.S. National Weather Service (http://www.cpc.ncep.noaa.gov).

11. Even though the Henry Hub is located in Louisiana, weather conditions in other consumption regions will affect the price at Henry Hub due to the integration of the national gas systems. We have selected representative cities for the weather index. Due to a lack of data, the weather variables are not volume weighted.

Table 2. Tests for Stationarity

Panel A. Augmented Dickey-Fuller Unit Root Tests

Variable	Lag	t_τ	Φ1	t_μ	Φ2	t*
Price	4	-3.88**				
Quantity	4	-3.562**				
Oil Price	4	-2.618	3.501	-2.263	2.576	-0.327
Storage	4	-4.434**				
Rig Count	4	-2.935	4.345	-1.396	1.128	0.047
IP	4	-1.003	1.59	-1.707	5.42**	
5% Critical Value		-3.41	6.25	-2.86	4.59	-1.95

Panel B. KPSS Tests for Stationarity

Variable	η_μ	η_τ	Lag	Trend Sig (p-value)
Price	1.68**	0.139	12	0.000
Quantity	1.248**	0.064	12	0.000
Oil Price	1.049**	0.380**	12	0.000
Storage	0.017	0.0168	12	0.761
Rig Count	3.294**	0.279**	12	0.000
IP	4.479**	0.661**	12	0.000
5% Critical Value	0.463	0.146		

** indicates significance at the 5% level.

The augmented Dickey-Fuller regression takes the following form: $\Delta X_t = \alpha_0 + \alpha_1 t + \alpha_2 X_{t-1} + \sum_{i=1}^{k} \alpha_{2+i} \Delta X_{t-i} + e_t$. t_τ tests the null hypothesis of unit root with trend, $\alpha_2 = 0$ while $\alpha_1 \neq 0$ (unit root with a trend), Φ1, the null of $\alpha_1 = 0$ while $\alpha_2 = 0$, t_μ the null of a unit root when a constant is present ($\alpha_2 = 0$ while $\alpha_0 \neq 0$), Φ2 the joint hypothesis of $\alpha_0 = 0$ while $\alpha_2 = 0$, (constant = 0 and root = 1), and t*, the null of a unit root without a constant or a trend. Results for a lag length of 4 are reported. The results are robust to various lag lengths.

The Dickey-Fuller test takes as the null hypothesis that a unit root is present in the series being examined. The null hypothesis of a unit root is rejected whenever the computed test statistic is less than the 5% critical value.

The KPSS test (Kwiatkowski, Phillips, Schmidt, and Shin 1992) takes as the null hypothesis that the series being examined is stationary. The test statistic η_μ is used to test null hypothesis of stationarity when a constant is present and η_τ to test the null of stationarity when a trend is present. The null is rejected whenever the test statistic's computed value exceeds the 5% critical value. A lag length of 12 is used following the usual suggestions in the literature.

unit root test (Dickey and Fuller 1979) and the KPSS test (Kwiatkowski, Phillips, Schmidt, and Shin 1992).[12]

Panel A of Table 2 presents the ADF test results. The results are robust to various lag lengths. The null hypothesis is rejected whenever the computed test statistic exceeds a critical value. We select the 5% level of significance as the benchmark and report the critical values in the last row of the table.

The test statistic t_τ leads to rejection of the null hypothesis of a unit root with trend for the series Price, Quantity, and Storage. Because this is the most stringent test, rejection of the null implies that we need look no further for these series but conclude they do not exhibit unit roots. On the other hand, the null is not rejected for the additional series. With this in mind several alternative hypotheses are tested. We do not reject the unit root null hypothesis for any of the remaining series for any alternative formulation of the test.

The KPSS test results, reported in Panel B of Table 2, corroborate the ADF test results. Accounting for a trend, the test does not reject stationarity for Price and Quantity. When a trend is accounted for in the tests of whether Oil Price, Rig Count and IP are stationary, we continue to reject stationarity. Stationarity of the variable Storage is never rejected, and it does not contain a significant trend. The results of the KPSS tests corroborate the inferences from the ADF tests. Therefore, we feel confident with the following inferences: Price and Quantity are trend-stationary; Oil Price, Rig Count, and IP each exhibit a trend and are non-stationary; Storage exhibits no trend and is stationary. Unreported results indicate that the first differences of Oil Price, Rig Count and IP are stationary.

D. A Further Look at Quantity

Inspection of the Quantity series revealed what appeared to be a structural shift in its level. Because a structural shift could influence our inferences about the stationarity of this variable, we formally tested the unit root hypothesis allowing for a potential structural shift. Two separate hypotheses are proposed and tested. IO1 (Innovative Outlier) posits a model in which only the intercept in the Dickey-Fuller regression experiences a shift. IO2 posits a model in which both the intercept shifts and there is a shift in the relation between Quantity and the time trend. The statistical tests are developed in Perron (1997). Specifically the formulations of the hypotheses are given by:

$$IO1: q_t = \mu_q + \theta DU_t + \beta t + \delta D(T_b)_t + \alpha q_{t-1} + \sum_{i=1}^{k} c_i \Delta q_{t-i} + e_t \tag{6}$$

$$IO2: q_t = \mu_q + \theta DU_t + \beta t + \gamma DT_t + \delta D(T_b)_t + \alpha q_{t-1} + \sum_{i=1}^{k} c_i \Delta q_{t-i} + e_t \tag{7}$$

12. See Greene (2000) and Hamilton (1994) for a complete development of the stationarity tests. All estimation and test results are computed using RATS/Regression Analysis of Time Series. The null for the Dickey-Fuller test is that the series is non-stationary (exhibits a unit root). The null for the KPSS test is that the series is stationary.

Table 3. Test for a Unit Root in the Variable Quantity Conditional on a Structural Break

Model	T_b	k	t_θ	t_δ	t_γ	α	t_α	Asymp. 5% c.v. of t_α
IO1	2000:4:21	1	3.047***	-1.846*		0.766	-4.918**	-4.80
IO2	2000:5:5	8	-1.194	-0.83	1.48	0.623	-4.829	-5.08

*, **, *** indicate significance at the 10, 5 and 1% levels, respectively.

The method follows a test in Perron (1997) for a unit root in a series with a structural break. The break is endogenously determined. The following innovative Outlier Models (IQ) are employed:

IO1: $\text{Quantity}_t = \mu + \theta DU_t + \beta t + \delta D(T_b)_t + \alpha Quantity_{t-1} + \sum_{i=1}^{k} c_i \, \Delta \, Quantity_{t-i} + e_t$,

IO2: $\text{Quantity}_t = \mu + \theta DU_t + \beta t + \gamma DT_t + \delta D(T_b)_t + \alpha Quantity_{t-1} + \sum_{i=1}^{k} c_i \, \Delta \, Quantity_{t-i} + e_t$,

where: $DU_t = 1$ for $t > T_b$, and T_b is the break point,
$D(T_b)_t = 1$ for $t = T_b + 1$,
$DT_t = 1*t$, for $t > T_b$, where t is the time trend.

where: q_t represents quantity at time t,
$DU_t = 1$ for $t > T_b$, and T_b is the break point,
$D(T_b)_t = 1$ for $t = T_b + 1$,
$DT_t = 1*t$, for $t > T_b$, and
t is the time trend variable.

The results are reported in Table 3. The t statistics for the tests that $\theta = 0$ and $\delta = 0$ in the IO1 model indicate significance for the intercept shift at the 1% level and significance at the 10% level for the innovation shock. The test statistic for the unit root hypothesis t_a is more negative than the critical 5% value, indicating that we reject the unit root hypothesis under the IO1 hypothesis. The results for the IO2 hypothesis do not support the proposition that there was a shift in the slope of the trend line. We conclude that IO1 better describes the variable Quantity, and that despite a shift in the intercept the variable does not exhibit a unit root. However, we also conclude that Quantity is trend stationary and exhibited a structural break. We account for the intercept shift by including a dummy variable in the models in which Quantity is an endogenous variable (DUMMY=1 if t > 2000:4:19 and 0 otherwise). We also conducted similar tests for a structural shift in the Price series. Those results (not reported) did not reject the null hypothesis that no structural shift had occurred.[13]

13. The results are available from the authors upon request.

Table 4. Univariate Time Series Model Estimation Results

Parameter	AR(4)	ARMA(3,1)	ARMAX
α_0	1.0097	0.9824	-2.155
	(0.066)	(0.0936)	(0.0001)
γ	0.0052	0.00526	0.005
	(0.002)	(0.003)	(0.0001)
α_1	1.3367	1.036	1.064
	(0.000)	(0.000)	(0.000)
α_2	-0.721	-0.326	
	(0.000)	(0.045)	
α_3	0.459	0.232	
	(0.000)	(0.000)	
α_4	-0.112		-0.224
	(0.0049)		(0.000)
ρ		0.295	
		(0.019)	
β_1			0.0013
			(0.945)
β_2			0.0034
			(0.000)
β_3			0.0053
			(0.303)
β_4			0.0374
			(0.000)
β_5			0.00009
			(0.958)
β_6			-0.0029
			(0.855)
β_7			0.00265
			(0.76)
β_8			-0.00246
			(0.006)
β_9			-22.40
			(0.114)

III. UNIVARIATE TIME-SERIES MODELS OF THE NATURAL GAS PRICE

Table 4 presents estimation results for three alternative univariate time-series models. The models are each identified using the AIC and BIC criteria from a search over ARMA and ARMAX model specifications. The Ljung-Box Q statistic for each model does not lead to rejection of the null hypothesis of white-noise errors.

The univariate models estimated take the following general form:

Table 4, continued. Univariate Time Series Model Estimation Results

Parameter	**AR(4)**	**ARMA(3,1)**	**ARMAX**
$\overline{R}^2$	0.9493	0.9492	0.983
AIC	2032.96	2034.46	255.51
BIC	2058.69	2060.19	292.06
Q-Stat	22.72 (Q(32))	24.18 (Q(32))	19.29 (Q(20))
Sig. Of Q	0.889	0.838	0.566
Sample Period	1991:2:1 - 2001:5:18	1991:2:1 - 2001:5:18	1999:1:1 - 2001:5:18

p-values are shown in parentheses
The univariate models estimated take the following general form:

$$p_t = \alpha_0 + \gamma t + \sum_{i=1}^{4} \alpha_i p_{t-i} + \sum_{k=1}^{K} \beta_k X_k + e_t, \quad e_t = \rho e_{t-1} + \varepsilon_t$$

Where the exogenous variables are X = ($\Delta OilPrice_{t-3}$, $Storage_{t-1}$, CDD_t, HDD_t, $\Delta RigCount_{t-1}$, ΔFY_CDD_{t-1}, ΔFY_HDD_{t-2}, $\Delta FY\_Storage_{t-1}$, ΔIP_t) and Δ denotes the first difference, *CDD* is the Cooling Degree Day series, *HDD* is the Heating Degree Day series, *IP* is the gas-weighted industrial output index, and ΔFY_X denotes the difference between the variable X and its normal value. The normal value for *CDD* and *HDD* is the 1970-2000 average and the normal value for storage is the previous 5-year average of injection or drawdown. AIC is the Akaike Information Criterion and BIC is the Schwarz Bayesian Information Criterion.

$$p_t = \alpha_0 + \gamma t + \sum_{i=1}^{4} \alpha_i p_{t-i} + \sum_{k=1}^{K} \beta_k X_k + e_t, \quad e_t = \rho e_{t-1} + \varepsilon_t \tag{8}$$

The variable p_t is the weekly spot price series and the variable t is the time trend. The X_k are the exogenous variables. The first two models for which results are reported in Table 4 are simple time-series models and both are presented because they are associated with virtually identical summary characteristics. The first model is an AR(4) model for Price and the second is an ARMA(3,1) model. Both models produce roughly equal adjusted R^2, AIC, and BIC values, and the Ljung-Box Q statistic for each does not lead to rejection of the null hypothesis that the errors are white noise.

Among other alternative models examined (not reported) are the MA(1) model and the AR(1) model. The latter model was in particular examined because of its affinity to the random walk for lag-1 correlations close to 1. In general, these alternative univariate models do not perform as well as those presented in the first two columns of Table 4. For example, the adjusted R-square for the MA(1) model is only 79.2%, and the AIC and BIC values are 2789.88 and 2802.75. The computed AIC and BIC values are much higher than those of the models reported. Since the price series is trend stationary, the random walk model is not really appropriate for constructing forecasts; instead, we settled on the AR(1) formulation. The AR(1) formulation is associated with an adjusted R-square of 93.9% and AIC and BIC values of 2128 and 2142, respectively. These results are inferior to the models presented. This of course is to be expected since the lag lengths for the models presented in Table 4 were chosen based upon optimizing the AIC and BIC criteria, along with assuring the residuals from the estimated models were not statistically different from white noise.

The third model is more refined in the sense that assumed exogenous variables are included along with lagged values of the price. Such a model is usually referred to as a member of the ARMAX family. The theory discussed earlier suggests that several potential factors may influence prices. Our point is not to test the theory directly but rather to let the theory guide us in terms of the inclusion of variables that can reasonably be expected to act as fundamental determinants of prices. We use cross-correlation statistics (not reported) between the natural gas price series and the exogenous variables of the system to identify candidate lag relations computed over the range t-6 through t+6. We find that, based upon the cross-correlation statistics, the following lag specifications have the highest correlation with the Price series

$$X' = \begin{pmatrix} \Delta Oil\ Price_{t-3}, Storage_{t-1}, CDD_t, HDD_t, \Delta Rig\ Count_{t-1}, \Delta FY\_CDD_{t-1}, \\ \Delta FY\_HDD_{t-2}, \Delta FY\_Storage_{t-1}, \Delta IP_t \end{pmatrix} \quad (9)$$

where Δ denotes the first difference and ΔFY_J denotes the difference between the variable J and its normal value. The normal values for *CDD* and *HDD* are the 1970–2000 averages and the normal value for Storage is the previous five-year average. We take first differences of Oil Price, Rig Count, and IP because our earlier results indicated that these variables are not stationary. The variable $\Delta Oil\ Price_{t-3}$ is equal to $\Delta Oil\ Price_{t-3} - \Delta Oil\ Price_{t-4}$. The weather variables are included for the reasons discussed earlier. In the general formulation of the univariate model shown in equation (8), the exogenous variables are represented as the X_k variables where k represents the position of the lag variable in the above vector X.

The estimation results for the ARMAX model are reported in column 3 of Table 4. The results are comparable to the univariate ARMA models on an adjusted R^2 basis. The AIC and BIC values for this model cannot be directly compared to the other models. The Q statistic for the errors of the ARMAX model does not reject white noise.[14]

IV. MULTIVARIATE MODELS

In this section we report the results from estimating multivariate models in which more than one of the variables is taken to be endogenous. Specifically, we first treat Price and Quantity as endogenous to the system and estimate two alternative specifications. The first model is a two-variable system that includes only Price and Quantity following the general *pth*-order vector autoregression

14. An alternative structure is the time-varying risk premium formulation of the expected spot price $E[S_{t+1}] = F_t + \lambda t$ where $E[S_{t+1}]$ is the expected spot price at t+1, F_t is the current futures price for delivery at t+1 and λ_t is the time-dependent risk premium (see for instance Pindyck 2001 and Chiou Wei and Zhu 2006). We do not explore this formulation in the current study but leave that for future research.

$$\begin{bmatrix} p_t \\ q_t \end{bmatrix} = \begin{bmatrix} \mu_p \\ \mu_q \end{bmatrix} + \begin{bmatrix} \gamma_p \\ \gamma_q \end{bmatrix} t + \begin{bmatrix} 0 \\ \beta_{qD} \end{bmatrix} D + \Gamma_1 \begin{bmatrix} p_{t-1} \\ q_{t-1} \end{bmatrix} + \cdots + \Gamma_p \begin{bmatrix} p_{t-p} \\ q_{t-p} \end{bmatrix} + \begin{bmatrix} \varepsilon_{p,t} \\ \varepsilon_{q,t} \end{bmatrix}$$

$$\underline{y_t} = \underline{\mu} + \underline{\gamma}t + \underline{\beta_D} D + \sum_{i=1}^{p} \left(\Gamma_i \underline{y_{t-i}} \right) + \underline{\varepsilon_t} \tag{10}$$

where underscores denote vectors, *t* is a time trend, *D* is a dummy taking the value 0 for each date from the beginning of the series through 2000:4:18 and 1 for the remaining dates, and Γ_t is a 2 x 2 matrix of coefficients for each lag 1 through *p*.[15] The coefficient vectors in (10) $\underline{y_t}, \underline{\mu}, \underline{\gamma}, \underline{\beta_D}, \underline{\varepsilon_t}$ are each 2 x 1 vectors, where 2 equals the number of endogenous variables. The general formulation of the model in which exogenous variables also appear is given by

$$\underline{y_t} = \underline{\mu} + \underline{\gamma}t + \underline{\beta_D} D + \sum_{i=1}^{p} \left(\Gamma_i \underline{y_{t-i}} + \Theta_i \underline{X_{t-i}} \right) + \underline{\varepsilon_t} \tag{11}$$

The model differs from equation (10) in the inclusion of the K x 1 vector *X* of exogenous variables where the notation $\underline{X_{t-i}}$ denotes the K-row vector of observations on exogenous variables measured at date *t-i* and Θ_i is a 2 x K matrix of coefficients.

Earlier we concluded that Price and Quantity are both trend-stationary variables. The two-variable system shown in (10) is therefore a standard bivariate VAR model where we account directly for the trend in estimation as well as the structural break identified for the quantity series. The number of lags is equal to 4 and is found by minimizing the AIC and BIC. Such a system is informative to the extent that it permits us to forecast one or the other endogenous variables, a point we will turn to in the next section. As a prelude to that analysis, we present statistics on the influence of shocks to variables considered to be endogenous on the endogenous variables of the system.

Panel A of Table 5 presents results for the bivariate VAR. The panel presents the percentage of forecast variance for Price and Quantity explained by either a shock to Price or a shock to Quantity where the time horizon is 20 weeks. The first column shows clearly that a shock to Price explains roughly 86.48% of the forecast variance in Price, while the second column shows that a shock to Quantity explains roughly 13.52% of the forecast variance in Price by week 20. Therefore, empirically in the North American natural gas market, Quantity does not play the significant role engendered by the model described in Section I. Accounting for the exogenous variables identified earlier has only a minor influence on the results as shown in the third and fourth columns of the panel.

Panel B of Table 5 extends the bivariate model by allowing Storage and ΔRig

15. It is well known that the VAR can be motivated by reference to the Wold Decomposition Theorem (Greene 2003).

table 5 as of 6-12-06 630 am.tif

Count to also enter as endogenous variables, while including the remaining variables as exogenous factors. There are N = 4 endogenous and K = 7 exogenous variables in the system. As in the bivariate VAR formulation, we also include a trend factor for those equations where the stationarity tests indicate it is required and control for the structural break in the Quantity series. The conclusions are roughly the same as those drawn from the results presented in Panel A. It is important to recognize however that in analyzing dynamics we do not shock the exogenous variables but only the endogenous variables in the system.

V. FORECASTING PERFORMANCE OF THE MODELS

A. Forecast Period 1: January 1, 2001–May 18, 2001

The utility of the models estimated in the prior two sections is their ability to forecast prices. The benchmark we use for assessing comparative ability is the natural gas futures price. In the next section we take up the question of the average cost of acquiring gas under various hedging strategies in comparison to a strategy of always hedging with NG futures. In this section we focus on the forecast performance of the Price forecasting models. We examine two price phases: Forecast Period 1 was a time of falling prices; in contrast Forecast Period 2, examined later, was a time of rising prices. Our approach is to hypothesize the existence of a decision maker who utilizes all the data available up through a given date τ when formulating price forecasts. We make the explicit assumption that the decision maker cannot predict whether prices are moving into a period of rising or falling phase. We feel this conforms with what might generally be expected in practice. Thus, for Forecast Period 1 the decision maker utilizes data commencing January 4, 1991, up through December 29, 2000, in the estimation of the models and constructs his forecasts based upon the estimated parameters. The data are then updated using the next week's data, the models are re-estimated, and new forecasts constructed.

The forecast performance results for the univariate and multivariate models, as well as for the futures price, are reported in Tables 6 and 7 for a period during which prices were generally declining. Results are presented for 3-, 6-, 9- and 12-month forecast horizons. The difference between the two tables lies in the assumption about the information available to the forecaster on the exogenous variables in the system. Table 6 represents the perfect foresight case and is meant to act as a benchmark. The statistics reported in Table 6 for the models including the exogenous variables utilize the actual values of the exogenous variables in the forecasts. Table 7 on the other hand substitutes what we call "normal" values for the actual exogenous variable values in the forecast.[16]

16. Specifically, the normal values are defined as follows. CDD and HDD equal the 30 year daily average from 1970-2000. Storage equals the five-year average injections/drawdowns of gas. The differences from normal values for CDD, HDD, and injection/drawdowns are all assumed to be zero. Changes in oil price, rig count, and IP variables are also assumed to be zero based upon the results in Table 2 showing that the levels of the oil price, rig count, and IP are nonstationary.

Table 6. Forecast Performance of the Models Using Actual Values for Exogenous Variables and the Forecast Performance of the Futures Price (Forecast Period 1: 2001/1/1 – 2001/5/18)

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var VAR-X	4-Var. VAR-X	Futures
3-Month							
MAE	1.563	1.667	1.27	1.982	1.934	1.41	1.46
MAE%	37.9%	40.7%	36.7%	50.9%	54%	39.5%	37%
RMSE	2.180	2.244	1.53	2.192	2.107	1.54	2.35
RMSE%	47%	48.9%	45.2%	54.5%	62%	45.1%	52%
Direction							
Correct	18	18	20	12	17	15	8
Wrong	2	2	0	8	3	5	6
Indeterminate	0	0	0	0	0	0	6
6-Month							
MAE	2.708	2.831	2.36	3.145	3.157	2.71	2.68
MAE%	100%	105%	92%	119%	123%	106%	104%
RMSE	3.654	3.754	2.39	3.435	3.186	2.73	7.42
RMSE%	127%	131%	96%	127%	128%	110%	225%
Direction							
Correct	18	18	20	13	18	11	7
Wrong	2	2	0	7	2	9	11
Indeterminate	0	0	0	0	0	0	2
9-Month							
MAE	3.111	3.238	2.92	3.418	3.901	3.29	3.14
MAE%	141%	147%	128%	152%	171%	144%	140%
RMSE	4.488	4.615	2.95	3.812	3.969	3.31	10.00
RMSE%	215%	220%	132%	175%	176%	147%	305%
Direction							
Correct	18	18	20	12	13	8	7
Wrong	2	2	0	8	7	12	13
Indeterminate	0	0	0	0	0	0	0

Table 6, continued.

12-Month							
MAE	2.768	2.899	2.11	2.878	3.383	2.72	2.25
MAE%	111%	116%	84%	113%	130%	105%	91%
RMSE	4.822	4.987	2.39	3.471	3.566	2.89	6.21
RMSE%	197%	204%	100%	143%	145%	119%	245%
Direction							
Correct	18	18	20	13	14	5	16
Wrong	2	2	0	7	6	15	0
Indeterminate	0	0	0	0	0	0	4

The models are estimated up to 2001:5:18, weekly on a rolling basis starting from 2001:1:1. The statistics presented are based on 20 forecast observations. The Mean Absolute Error (MAE) is calculated as

$$MAE = \frac{\sum_{i=1}^{20} \left| Forecasted\ Price_i - Actual\ Price_i \right|}{20}$$

. The Root Mean Square Error (RMSE) is calculated as

$$RMSE = \left\{ \frac{\sum_{i=1}^{20} \left(Forecasted\ Price_i - Actual\ Price_i \right)^2}{20} \right\}^{.5}$$

We compute what we call the percentage error as

%Error = [(Forecast - Actual)/Actual]/100 of the week. We then calculate the absolute % error and then the average absolute % error. We label the result MAE%. We compute a similar measure using the RMSE and label that variable RMSE%. Forecast performance in terms of predicted direction is also tabulated. Define $p^F_{t+M} - p_t = \delta_{t+M}$ as the difference between the M-month ahead predicted price and today's actual price. Likewise define $p_{t+M} - p_t = \delta_{t+M}$ as the difference between the actual M-month ahead price and today's price. If sign (δ^F_{t+M}) = sign (δ_{t+M}) then we consider the forecast performance of the model to be "Correct" in terms of predicting the sign of the change in price relative to today's price. If sign (δ^F_{t+M}) $\neq$ sign (δ_{t+M}) and $\left|\delta^F_{t+M}\right| - \left|\delta_{t+M}\right| > .10$ then we classify the prediction as "Wrong." If sign (δ^F_{t+M}) $\neq$ sign (δ_{t+M}) and $\left|\delta^F_{t+M}\right| - \left|\delta_{t+M}\right| < .10$ then we classify the case as indeterminate.

Table 7. Forecast Performance of the Models Using Normal Values for Exogenous Variables and the Forecast Performance of the Futures Price
Forecast Period 1: 2001/1/1 – 2001/5/18

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var VAR-X	4-Var. VAR-X	Futures
3-Month							
MAE	1.563	1.667	0.94	1.982	1.72	1.41	1.46
MAE%	37.9%	40.7%	29.1%	50.9%	48%	39%	37%
RMSE	2.180	2.244	1.14	2.192	1.97	1.52	2.35
RMSE%	47%	48.9%	37%	54.5%	56%	44%	52%
Direction							
Correct	18	18	17	12	17	15	8
Wrong	2	2	3	8	3	5	6
Indeterminate	0	0	0	0	0	0	6
6-Month							
MAE	2.708	2.831	1.93	3.146	2.98	2.69	2.68
MAE%	100%	105%	81%	119%	114%	105%	104%
RMSE	3.654	3.754	2.03	3.436	3.03	2.72	7.42
RMSE%	127%	131%	87%	127%	121%	109%	225%
Direction							
Correct	18	18	17	13	17	12	7
Wrong	2	2	3	7	3	8	11
Indeterminate	0	0	0	0	0	0	2

The results and statistics reported are obtained in the same fashion as in Table 6. The difference in these computations is that the forecasts are based on normal values for exogenous variables. Specifically, CDD and HDD equal the 30 year daily average from 1970-2000. Storage equals the five-year average injections/drawdowns of gas. The differences from normal values for CDD, HDD and injection/drawdowns are all assumed to be zero. Changes in oil price, rig count, and IP variables are also assumed to be zero based upon the results in Table 2 showing that the levels of the oil price, rig count and IP are nonstationary. For ease of comparison, the results of AR(4), ARMA(3,1) and 2-variable VAR models are retained in this table.

Table 7, continued. Forecast Performance of the Models Using Normal Values for Exogenous Variables and the Forecast Performance of the Futures Price

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var VAR-X	4-Var. VAR-X	Futures
9-Month							
MAE	3.111	3.238	3.52	3.418	4.37	3.67	3.14
MAE%	141%	147%	159%	152%	186%	161%	140%
RMSE	4.488	4.615	3.59	3.812	4.54	3.69	10.00
RMSE%	215%	220%	166%	175%	196%	165%	305%
Direction							
Correct	18	18	15	13	8	6	7
Wrong	2	2	5	7	12	14	13
Indeterminate	0	0	0	0	0	0	0
12-Month							
MAE	2.768	2.899	2.67	3.418	3.90	3.11	2.25
MAE%	111%	116%	109%	113%	149%	121%	91%
RMSE	4.822	4.987	2.96	2.877	4.12	3.33	6.21
RMSE%	197%	204%	127%	143%	169%	137%	245%
Direction							
Correct	18	18	20	13	11	4	16
Wrong	2	2	0	7	9	16	0
Indeterminate	0	0	0	0	0	0	4

The results and statistics reported are obtained in the same fashion as in Table 6. The difference in these computations is that the forecasts are based on normal values for exogenous variables. Specifically, CDD and HDD equal the 30 year daily average from 1970-2000. Storage equals the five-year average injections/drawdowns of gas. The differences from normal values for CDD, HDD and injection/drawdowns are all assumed to be zero. Changes in oil price, rig count, and IP variables are also assumed to be zero based upon the results in Table 2 showing that the levels of the oil price, rig count and IP are nonstationary. For ease of comparison, the results of AR(4), ARMA(3,1) and 2-variable VAR models are retained in this table.

Two statistics are computed as measures of the accuracy of the forecasts. MAE is the mean absolute deviation of the forecast price from the actual price and is computed as

$$MAE = \frac{\sum_{i=1}^{20} \left| Forecasted\ Price_i - Actual\ Price_i \right|}{20} \tag{6}$$

The number of forecasts possible given the length of our time series is equal to 20. The Root Mean Squared Error (RMSE) is computed as

$$RMSE = \left\{ \frac{\sum_{i=1}^{20} \left(Forecasted\ Price_i - Actual\ Price_i \right)^2}{20} \right\}^{.5} \tag{7}$$

The initial time series used to estimate a model is defined generically as t through T1. The forecast is then for the time horizon T1 + M. The model is then re-estimated using data for the period t through T1 +1, and a new forecast is computed for T1 + 1 + M. That is, the forecasts are constructed on a rolling basis. In addition, the forecast is dynamic, that is, the forecasted value of the price for an earlier period is fed into the forecast process to generate price forecasts for a later period due to the fact that the price is modeled as an autoregressive process. This procedure allows us to include as many forecast periods as is legitimately possible.

Tables 6 and 7 also present statistics on the relative percentage error. We compute what we call the percentage error as %Error = [(Forecast - Actual)/Actual]/100 of the week. We then calculate the absolute % error and then the average absolute % error. We label the result MAE%. We compute a similar measure using the RMSE and label that variable RMSE%. Finally, we present the number of times the model predicts the direction of the price move correctly relative to the forecast date *t*. These statistics are found by first computing the actual sign of the change in price between today's price and the actual price at the forecast horizon date. The sign of the predicted change is then measured as the difference between the forecasted price and today's price. The count measures assess how many times the model predicted the correct sign of the change.

The models that include the exogenous variables tend to produce better predictions in terms of MAE and RMSE. The ARMAX model performs the best not only when compared to the other statistical models but also when compared with the futures price. The ARMAX model performs the best whether we treat the forecaster as having perfect knowledge of the contemporaneous variables (Table 6), or whether the forecaster is assumed to use data only observable at the date the forecast is constructed (Table 7). Further, the ARMAX model is best at all forecast horizons. The ARMAX model also beats the futures price. In terms of the percentage error relative to the actual price, the model dominates all others. However, the average absolute percentage errors are relatively large and suggest the model works best over the short three-month forecast horizon. Nevertheless, the model predicts the direction of change in the price 100% of the time. We conclude from these results that forecast power is maximized using a simple ARMAX time series model of the spot price during the price regime examined.

We now turn to the issue of the cost of gas under three alternative hedging strategies. We present values of the average cost of natural gas to the user who either locks in the price with a futures contract (hedges) or transacts on the spot market. The choice is assumed to depend upon a forecast of the spot price. Natural gas price forecasts are obtained as in Table 6. Similar conclusions on the relative merits of each forecasting model occur if we follow the information assumptions of Table 7. The cost of the gas if the choice is to always hedge is the M-month ahead futures price at the time of the forecast. That is, if the choice is to always hedge, the user buys a futures contract for the full planned requirement. We also propose a mixed strategy that sometimes involves buying a futures contract and sometimes involves no hedging. The choice to hedge under the mixed strategy is based on a comparison of the price forecast and the futures price for the period M-months ahead. The cost of the gas under the mixed hedging strategy is the futures price at the time of the forecast when the M-month ahead futures price is lower than the M-month ahead price forecast. Finally, we present the cost to the user of never hedging. The cost of the gas when there is no hedging is the actual price that prevailed at the end of the forecast horizon.

For convenience we assume that the gas requirement is always one unit, and for each week, a decision of whether to hedge the one-unit gas requirement is made for different horizons (3-, 6-, 9-, and 12-month). Futures prices for the dates required were obtained from NYMEX. The futures contract traded on the NYMEX is for gas delivered at the Henry Hub. Since our models are based on spot prices at Henry Hub, the futures prices and the spot prices are for comparable gas.

Panel A of Table 8 illustrates the computation of the average cost of the Always Hedge strategy, the Mixed Strategy and the No Hedge strategy. The price-forecasting model used in Panel A is the univariate AR(4) specification and the time horizon is three months. In this particular case the Mixed Strategy produces a lower average cost ($4.12) than the Always Hedge strategy $5.39. However, the No Hedge strategy produces the lowest average cost of gas, $3.93.

Panel B presents the average cost numbers for each forecasting model, each strategy and each forecast time horizon. The Always Hedge values are constant across each model since the futures price is not a function of the strategy or the price-forecasting model. The No Hedge average cost is also independent of the strategies and the models since the cost of gas is based on the actual gas prices for the delivery month. The forecasting model associated with the minimum average cost for the Mixed Strategy is the same for each forecast horizon, and as would be expected from the results in Tables 6 and 7, it is the univariate ARMAX model. In

Table 8. Average Cost of Gas Under the Always Hedge, Never Hedge and Mixed Strategies (Forecast Period 1: 2001/1/1 – 2001/5/18)

A. Sample Calculation of Cost of Gas - AR(4) Model for 3-month Ahead Forecast

Forecast Week	3-Month Ahead Forecast for Week	3-Month Ahead Price Forecast	3-Month Ahead Futures Price	Actual Price 3 Months After Forecast	Cost of Gas Always Hedge	Cost of Gas Mixed Strategy
1/5/01	4/6/01	11.30	6.26	5.25	6.26	6.26
1/12/01	4/13/01	11.15	6.54	5.41	6.54	6.54
1/19/01	4/20/01	7.47	6.01	5.25	6.01	6.01
1/26/01	4/27/01	6.79	6.06	4.99	6.06	6.06
2/2/01	5/4/01	5.23	5.57	4.63	5.57	4.63
2/9/01	5/11/01	5.15	5.70	4.30	5.70	4.30
2/16/01	5/18/01	4.77	5.50	4.29	5.50	4.29
2/23/01	5/25/01	4.52	5.19	4.08	5.19	4.08
3/2/01	6/1/01	4.46	5.35	3.78	5.35	3.78
3/9/01	6/8/01	4.53	5.18	3.77	5.18	3.77
3/16/01	6/15/01	4.38	5.09	3.86	5.09	3.86
3/23/01	6/22/01	4.52	5.38	3.82	5.38	3.82
3/30/01	6/29/01	4.64	5.33	3.46	5.33	3.46
4/6/01	7/6/01	4.56	5.48	3.00	5.48	3.00
4/13/01	7/13/01	4.76	5.49	3.13	5.49	3.13
4/20/01	7/20/01	4.57	5.24	3.09	5.24	3.09
4/27/01	7/27/01	4.49	4.94	3.04	4.94	3.04
5/4/01	8/3/01	4.26	4.64	3.17	4.64	3.17
5/11/01	8/10/01	4.09	4.42	3.07	4.42	3.07
5/18/01	8/17/01	4.12	4.45	3.11	4.45	3.11
Average Price				**3.93**	**5.39**	**4.12**

The forecasts are obtained as in Table 6. For the Always Hedge case, the cost of the gas is always the M-month ahead forward price at the time of the forecast. For the Mixed Strategy, hedging is selected based on a comparison of the price forecast and the futures price for the period M-months ahead. The cost of the gas is the futures price at the time of the forecast when the M-month ahead futures price is lower than the M-month ahead price forecast. Otherwise, there is no hedging. The cost of gas when there is no hedging is the actual price for the forecast month.

Table 8, continued. Average Cost of Gas Under the Always Hedge, Never Hedge and Mixed Strategies (Forecast Period 1: 2001/1/1 – 2001/5/18)

B. Cost of Gas Comparison (in $/MMBTU

Forecast Horizon and Hedge Type	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X
3-month						
No Hedge	3.93	3.93	3.93	3.93	3.93	3.93
Always Hedge	5.39	5.39	5.39	5.39	5.39	5.39
Mixed Strategy	4.12	4.12	4.03	4.60	4.18	4.33
6-month						
No Hedge	2.67	2.67	2.67	2.67	2.67	2.67
Always Hedge	5.36	5.36	5.36	5.36	5.36	5.36
Mixed Strategy	3.20	3.20	2.80	4.01	3.01	3.77
9-month						
No Hedge	2.34	2.34	2.34	2.34	2.34	2.34
Always Hedge	5.49	5.49	5.49	5.49	5.49	5.49
Mixed Strategy	3.02	3.02	2.74	3.54	4.39	5.34
12-month						
No Hedge	2.86	2.86	2.86	2.86	2.86	2.86
Always Hedge	5.11	5.11	5.11	5.11	5.11	5.11
Mixed Strategy	3.39	3.39	3.13	4.63	5.11	5.11

The forecasts are obtained as in Table 6. For the Always Hedge case, the cost of the gas is always the M-month ahead forward price at the time of the forecast. For the Mixed Strategy, hedging is selected based on a comparison of the price forecast and the futures price for the period M-months ahead. The cost of the gas is the futures price at the time of the forecast when the M-month ahead futures price is lower than the M-month ahead price forecast. Otherwise, there is no hedging. The cost of gas when there is no hedging is the actual price for the forecast month.

most cases, the Mixed Strategy produces a lower average cost than the Always Hedge strategy, irrespective of the forecasting model being employed. However, for the time period examined, the No Hedge strategy would have produced the lowest average cost for each time horizon.

B. Forecast Period 2: July 1, 2000 - December 31, 2000

Forecast Period 2 was a time of rising prices. Table 9 presents forecast performance results for Period 2. Table 9 is presented in the same format as Table 6; however, in contrast to Table 6, the futures price tends to dominate in terms of MAE and RMSE. Perhaps most importantly during this price phase, the percentage forecast error (MAE% or RMSE%) is smallest when the futures price is used as the predictor out to six months but performs worse than the four-variate VAR model that includes exogenous variables for longer horizons.

The same is true for the results presented in Table 10, which mirror those presented in Table 9. Table 11 presents the average cost of gas for the three strategies.

Table 9. Forecast Performance of the Models Using Actual Values for Exogenous Variables and the Forecast Performances of the Futures Price (Forecast Period 2: 2000/7/1 – 2000/12/31)

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X	Futures
3-month							
MAE	2.809	2.814	2.59	2.196	2.59	1.83	1.06
MAE%	41%	41%	38%	31.4%	39%	22%	13%
RMSE	3.257	3.263	3.16	2.659	3.380	2.11	1.28
RMSE%	45%	45.1%	46.5%	35.2%	54%	26%	15%
Direction							
Correct	4	4	5	9	9	24	15
Wrong	21	21	20	16	16	1	8
Indeterminate	0	0	0	0	0	0	1
6-month							
MAE	2.544	2.547	1.87	1.989	2.000	1.37	1.17
MAE%	45%	45%	33.6%	36.5%	38%	19%	14%
RMSE	3.211	3.216	2.52	2.605	3.050	1.76	2.38
RMSE%	57%	57%	45.3%	47.8%	66%	26%	38%
Direction							
Correct	7	7	11	8	13	19	10
Wrong	18	18	14	17	12	6	15
Indeterminate	0	0	0	0	0	0	0

The models are estimated up to 2001:5:18, weekly on a rolling basis starting from 2001:1:1. The statistics presented are based on 20 forecast observations.
The Mean Absolute Error (MAE) is calculated as

$$MAE = \frac{\sum_{i=1}^{20} \left| Forecasted\ Price_i - Actual\ Price_i \right|}{20}.$$

The Root Mean Square Error (RMSE) is calculated as

$$RMSE = \left\{ \frac{\sum_{i=1}^{20} \left(Forecasted\ Price_i - Actual\ Price_i \right)^2}{20} \right\}^{.5}.$$

We compute what we call the percentage error as %Error = [(Forecast - Actual)/Actual]/100 of the week. We then calculate the absolute % error and then the average absolute % error. We label the result MAE%. We compute a similar measure using the RMSE and label that variable RMSE%. Forecast performance in terms of predicted direction is also tabulated. Define $p_{t+M}^F - p_t = \delta_{t+M}$ as the difference between the M-month ahead predicted price and today's actual price. Likewise define $p_{t+M} - p_t = \delta_{t+M}$ as the difference between the actual M-month ahead price and today's price. If sign (δ_{t+M}^F) = sign (δ_{t+M}) then we consider the forecast performance of the model to be "Correct" in terms of predicting the sign of the change in price relative to today's price. If sign (δ_{t+M}^F) $\neq$ sign (δ_{t+M}) and $\left|\delta_{t+M}^F\right| - \left|\delta_{t+M}\right| > .10$ then we classify the prediction as "Wrong." If sign (δ_{t+M}^F) $\neq$ sign (δ_{t+M}) and $\left|\delta_{t+M}^F\right| - \left|\delta_{t+M}\right| < .10$ then we classify the case as indeterminate.

Table 9, continued. Forecast Performance of the Models Using Actual Values for Exogenous Variables and the Forecast Performances of the Futures Price (Forecast Period 2: 2000/7/1 – 2000/12/31)

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X	Futures
9-month							
MAE	1.607	1.606	1.96	1.829	2.34	1.87	1.06
MAE%	54%	53%	67.2%	64.2%	83%	48%	25%
RMSE	2.550	2.562	2.38	2.489	3.29	1.82	12.64
RMSE%	105%	105%	92%	98.9%	135%	65%	115%
Direction							
Correct	19	19	10	14	8	23	19
Wrong	6	6	15	11	17	2	6
Indeterminate	0	0	0	0	0	0	0
12-month							
MAE	1.439	1.439	2.95	2.553	3.42	3.28	1.05
MAE%	63%	63%	124%	109%	143%	97%	43%
RMSE	2.596	2.615	3.11	3.022	3.83	2.67	3.09
RMSE%	108%	108%	137%	137%	164%	108%	130%
Direction							
Correct	24	24	5	14	3	18	24
Wrong	1	1	20	11	22	7	0
Indeterminate	0	0	0	0	0	0	1

The models are estimated up to 2001:5:18, weekly on a rolling basis starting from 2001:1:1. The statistics presented are based on 20 forecast observations.
The Mean Absolute Error (MAE) is calculated as

$$MAE = \frac{\sum_{i=1}^{20} \left| Forecasted\ Price_i - Actual\ Price_i \right|}{20}.$$

The Root Mean Square Error (RMSE) is calculated as

$$RMSE = \left\{ \frac{\sum_{i=1}^{20} \left(Forecasted\ Price_i - Actual\ Price_i \right)^2}{20} \right\}^{.5}.$$

We compute what we call the percentage error as %Error = [(Forecast - Actual)/Actual]/100 of the week. We then calculate the absolute % error and then the average absolute % error. We label the result MAE%. We compute a similar measure using the RMSE and label that variable RMSE%. Forecast performance in terms of predicted direction is also tabulated. Define $p^F_{t+M} - p_t = \delta_{t+M}$ as the difference between the M-month ahead predicted price and today's actual price. Likewise define $p_{t+M} - p_t = \delta_{t+M}$ as the difference between the actual M-month ahead price and today's price. If sign (δ^F_{t+M}) = sign (δ_{t+M}) then we consider the forecast performance of the model to be "Correct" in terms of predicting the sign of the change in price relative to today's price. If sign (δ^F_{t+M}) ≠ sign (δ_{t+M}) and $\left|\delta^F_{t+M}\right| - \left|\delta_{t+M}\right| > .10$ then we classify the prediction as "Wrong." If sign (δ^F_{t+M}) ≠ sign (δ_{t+M}) and $\left|\delta^F_{t+M}\right| - \left|\delta_{t+M}\right| < .10$, then we classify the case as indeterminate.

Table 10. Forecast Performance of the Models Using Normal Values for Exogenous Variables and the Forecast Performances of the Futures Price (Forecast Period 2: 2000/7/1 – 2000/12/31)

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X	Futures
3-month							
MAE	2.809	2.814	2.856	2.196	2.53	1.51	1.06
MAE%	41%	41%	44%	31.4%	37.6%	21%	13%
RMSE	3.257	3.263	3.1496	2.659	3.240	2.89	1.28
RMSE%	45%	45.1%	46.5%	35.2%	50%	25.1%	15%
Direction							
Correct	4	4	4	9	8	24	15
Wrong	21	21	21	16	17	1	8
Indeterminate	0	0	0	0	0	0	1
6-month							
MAE	2.544	2.547	2.267	1.989	1.940	1.092	1.17
MAE%	45%	45%	45%	36.5%	36.3%	18%	14%
RMSE	3.211	3.212	3.038	2.605	2.946	1.758	2.38
RMSE%	57%	57%	52%	47.8%	62.7%	25%	38%
Direction							
Correct	7	7	8	8	13	21	10
Wrong	18	18	17	17	12	4	15
Indeterminate	0	0	0	0	0	0	0

The results and statistics reported are obtained in the same fashion as in Table 9. The difference in these computations is that the forecasts are based on normal values for exogenous variables. Specifically, CDD and HDD equal the 30 year daily average from 1970-2000. Storage equals the five-year average injections/drawdowns of gas. The differences from normal values for CDD, HDD and injection/drawdowns are all assumed to be zero. Changes in oil price, rig count, and IP variables are also assumed to be zero based upon the results in Table 2 showing that the levels of the oil price, rig count, and IP are nonstationary. For ease of comparison, the results of AR(4), ARMA(3,1), and 2-variable VAR models are retained in this table.

We now see that the Mixed Strategy in which the prediction model is the four-variate VAR model with exogenous variables dominates for the three-month horizon and is equivalent to the Always Hedge Strategy for the six-month horizon. The Mixed Strategy always produces an average cost that is less than or equal to the Always Hedge Strategy for the 9 and 12 month horizons. However, the No Hedge strategy produces the least cost for the 9 and 12 month horizons.

At this point it is useful to consider a potential explanation for the size of the percentage forecast errors reported in Tables 6, 7, 9, and 10. During the two periods we examine, natural gas prices experienced very large price changes. The period covering January through May 2001 was associated with a huge price swing (see Figure 1). The average price during January 2001 was $8.23 and the average price during May 2001 was $4.17. Thus, forecasts predicated on historic time series data especially at the 6- and 12-month horizon are likely to produce large errors. Nevertheless, as Table 6 shows, the ARMAX model predicted the direction of the price change with perfect accuracy. The period from July 2000 to December 2000

Table 10, continued. Forecast Performance of the Models Using Normal Values for Exogenous Variables and the Forecast Performances of the Futures Price (Forecast Period 1: 2000/7/1 – 2000/12/31)

Model	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X	Futures
9-month							
MAE	1.607	1.606	1.93	1.829	2.290	1.52	1.06
MAE%	54%	53%	74%	64.2%	81.8%	48%	25%
RMSE	2.550	2.562	2.674	2.489	3.290	1.75	2.64
RMSE%	105%	105%	119%	98.9%	136%	66%	115%
Direction							
Correct	19	19	22	14	9	20	19
Wrong	6	6		11	16	5	6
Indeterminate	0	0	0	0	0	0	0
12-month							
MAE	1.439	1.439	2.767	2.553	3.68	2.54	1.05
MAE%	63%	63%	116%	109%	154%	106%	43%
RMSE	2.596	2.614	3.327	3.021	4.40	2.65	3.09
RMSE%	108%	108%	141%	137%	188%	118%	130%
Direction							
Correct	24	24	16	14	1	13	24
Wrong	1	1	9	11	24	12	0
Indeterminate	0	0	0	0	0	0	1

The results and statistics reported are obtained in the same fashion as in Table 9. The difference in these computations is that the forecasts are based on normal values for exogenous variables. Specifically, CDD and HDD equal the 30 year daily average from 1970-2000. Storage equals the five-year average injections/drawdowns of gas. The differences from normal values for CDD, HDD and injection/drawdowns are all assumed to be zero. Changes in oil price, rig count, and IP variables are also assumed to be zero based upon the results in Table 2 showing that the levels of the oil price, rig count, and IP are nonstationary. For ease of comparison, the results of AR(4), ARMA(3,1), and 2-variable VAR models are retained in this table.

is also associated with a huge price swing (see Figure 1). The average price during July 2000 was \$3.97 while the average price during December 2000 was \$8.95. Thus, the same observation applies.

VII. SUMMARY AND CONCLUSIONS

The market we examine is for natural gas traded through the Henry Hub at Henry, Louisiana. Futures contracts for gas deliverable at this location are actively traded on the NYMEX. The study begins by focusing on the forecasting precision associated with univariate as well as multivariate models involving the spot price of natural gas. Some of the models include only endogenous variables while other models also include fundamental exogenous variables predicted to influence prices by the economic theory of price formation for a storable commodity. We examine two distinct periods, one during which NG prices were falling and another during which prices were rising. We find that a univariate time series model that

Table 11. Average Cost of Gas Under the Always Hedge, Never Hedge, and Mixed Strategies (Forecast Period 2: 2001/7/1 – 2001/12/31)

Cost of Gas Comparison (in $/MMBTU

Forecast Horizon and Hedge Type	AR(4)	ARMA (3,1)	ARMA -X	2-Var. VAR	2-Var. VAR-X	4-Var. VAR-X
3-month						
No Hedge	6.47	6.47	6.47	6.47	6.47	6.47
Always Hedge	5.26	5.26	5.30	5.26	5.26	5.26
Mixed Strategy	6.58	6.58	6.64	6.40	5.69	5.05
6-month						
No Hedge	5.50	5.50	5.50	5.50	5.50	5.50
Always Hedge	4.53	4.53	4.53	4.53	4.53	4.53
Mixed Strategy	5.66	5.66	5.66	5.10	4.59	4.53
9-month						
No Hedge	3.64	3.64	3.64	3.64	3.64	3.64
Always Hedge	4.28	4.28	4.28	4.28	4.28	4.28
Mixed Strategy	3.97	3.97	3.97	4.33	4.27	4.27
12-month						
No Hedge	2.57	2.57	2.57	2.57	2.57	2.57
Always Hedge	4.25	4.25	4.25	4.25	4.25	4.25
Mixed Strategy	2.93	2.93	2.93	3.95	4.25	4.25

The forecasts are obtained as in Table 9. For the Always Hedge case, the cost of the gas is always the M-month ahead forward price at the time of the forecast. For the Mixed Strategy, hedging is selected based on a comparison of the price forecast and the futures price for the period M-months ahead. The cost of the gas is the futures price at the time of the forecast when the M-month ahead futures price is lower than the M-month ahead price forecast. Otherwise, there is no hedging. The cost of the when there is no hedging is the actual price for the forecast month.

incorporates fundamental variables related to production, storage, weather, and aggregate output performs best in a root mean square error sense among all the models examined when NG prices are falling. In contrast when prices are rising, a VAR specification with multiple fundamental endogenous and exogenous variables gives the best predictions for time horizons of either 6, 9 or 12 months, while the futures price gives the best predictions for a three-month horizon.

We also examine the average gas cost to a user who implemented either an Always Hedge, Never Hedge, or a Mixed Strategy based upon strategically selecting to hedge or not to hedge based upon the price forecast. The Mixed Strategy utilizes the spot price forecasts based upon the alternative forecasting models. Several different forecast time horizons are examined. We find that during the falling price phase, but irrespective of the forecast model utilized, the Mixed Strategy always produces an average cost that is less than or equal to the strategy of always hedging. The same is true during a rising price phase for the 9- and 12-month time horizons, but during the 3- and 6-month time horizons the policy of always hedging dominates. However, we also find that during a falling price phase the absolute least cost strategy is to never hedge, and that this is also the dominating strategy during the rising price phase for 9- and 12-month out horizons.

Our results have important implications for natural gas users. Specifically, the results suggest that strategically hedging based upon price forecasts can be an optimal strategy if the phase of the price cycle is not entirely clear.

References

Bresnahan, T.F. and Suslow, V.Y., 1985, Inventories as an Asset: The Volatility Of Copper Prices. *International Economic Review*, 26, 409-424.

Brown, G.W., Crabb, P.R., and Haushalter, D., 2005, Are Firms Successful At Selectively Hedging? *Working paper, University of North Carolina.*

Chambers, M. and Bailey, R., 1996, A Theory Of Commodity Price Fluctuations. *Journal of Political Economy*, 104, 924-957.

Chiou Wei, S.Z. and Zhu, Z., forthcoming, Commodity Convenience Yield and Risk Premium Determination: The Case of the U.S. Natural Gas Market, *Energy Economics*.

Clewlow, L. and Strickland, C., 2000, *Energy Derivatives: Pricing and Risk Management* (Lacima Publications, London).

Deaton, A. and Laroque, G., 1992, On The Behavior Of Commodity Prices. *Review of Economic Studies*, 59, 1-23.

Deaton, A. and Laroque, G., 1996, Competitive Storage And Commodity Price Dynamics. *Journal of Political Economy*, 104, 826-923.

Dickey, D.A. and Fuller, W.A., 1979, Distribution Of The Estimators for Autoregressive Time Series With A Unit Root. *Journal of American Statistical Association*, 74, 427-431.

Eydeland, A. and Wolyniec, K., 2003, *Energy and Power Risk Management* (John Wiley & Sons, Hoboken, NJ).

Fitzgerald, J. and Pokalsky, J.T., 1995, The Natural Gas Market. In *Managing Energy Price Risk* (Risk Publications, London).

Greene, W.H., 2000, *Econometric Analysis*, 4th ed. (Prentice-Hall, Upper Saddle River, NJ).

Gustafson, R.L., 1958, *Carryover Levels for Grains* (USDA, Washington, DC).

Hamilton, J., 1994, *Time Series Analysis* (Princeton University Press, Princeton, NJ).

Jarque, C.M. and Bera, A.K., 1987, A Test for Normaility of Observations and Regression Residuals. *International Statistical Review*, 55, 163-172.

Knill, A.M., Minnick, K., and Nejadmalayeri, A., forthcoming , Selective Hedging, Information Asymmetry and Futures Prices, *Journal of Business*.

Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., and Shin, Y., 1992, Testing The Null Hypothesis Of Stationarity Against The Alternative Of A Unit Root: How Sure Are We That Economic Time Series Have A Unit Root? *Journal of Econometrics*, 54, 159-178.

Leong, K.S., 1997, The Forward Curve In The Electricity Market. In *The U.S. Power Market* (Risk Publications, London).

Murry, D. and Zhu, Z., 2004, Enrononline And Informational Efficiency in the U.S. Natural Gas Market. *The Energy Journal*, 25, 57-74.

Muth, J.F., 1961, Rational Expectations and the Theory of Price Movements. *Econometrica*, 29, 315-335.

Perron, P., 1997, Further Evidence On Breaking Trend Functions In Macroeconomic Variables. *Journal of Econometrics*, 80, 355-385.

Pindyck, R., 2004, Volatility and Commodity Price Dynamics. *Journal of Futures Markets*, 24, 1029-1047.

Samuelson, P., 1971, Stochastic Speculative Price. *Proceedings of the National Academy of Sciences*, 68, 335-337.

Schwartz, E., 1997, The Stochastic Behavior Of Commodity Prices, Implications For Valuation and Hedging (Presidential Address). *Journal of Finance*, 52, 923-973.

Stulz, R., 1996, Rethinking Risk Management. *Journal of Applied Corporate Finance*, 9, 8-24.

Sturm, F.J., 1997, *Trading Natural Gas* (Penn Well Books, Tulsa, OK).

Williams, J.C. and Wright, B.D., 1991, *Storage and Commodity Markets* (Cambridge University Press, Cambridge, MA).