

# An Analysis of Price Volatility in Natural Gas Markets


This article presents an analysis of price volatility in the spot natural gas market, with particular emphasis on the Henry Hub in Louisiana. The purpose is to address whether natural gas prices have been more volatile in recent years and identify potential market factors that may contribute to price volatility. In addition to a first-order autoregressive error model, several graphical and statistical tools are used to examine trends and determine influencing factors. Although there is no demonstrated long-term trend in volatility, there are seasonal patterns and volatility is correlated strongly with storage dynamics. This report was written by Erin Mastrangelo. Questions or comments should be directed to William Trapmann at [william.trapmann@eia.doe.gov](mailto:william.trapmann@eia.doe.gov) or (202)586-6408.

## Introduction

The subject of price volatility in natural gas markets has received increased attention in recent years as the market experienced expanding dips and swells in prices while overall prices shifted to a higher level (Figure 1). Volatility is not defined by the level of prices, however, but by the degree of price variation in the market. Therefore, increasing natural gas prices do not necessarily indicate whether a market is volatile. Given that volatility is measured by percent differences in the day-to-day price

of natural gas, a large price movement at higher prices may equate to a comparable level of volatility as a smaller price movement when natural gas prices are lower. Although volatility is a key measure of natural gas market movements and fundamentals, expanding daily price movements at any volatility level can have vast impacts on traders and consumers of natural gas. When addressing price risk, it can be important to examine absolute price movements as well as volatility.

Figure 1: Daily Henry Hub Spot Price, 1994-2006



Source: NGI's Daily Gas Price Index, Intelligence Press

Market prices respond to shifts in supply and demand, and the degree of price response relates to the price elasticity of both. Natural gas prices have been particularly sensitive to short-term supply and demand shifts in recent years because of the highly inelastic nature of this market. In the short-term, consumers are limited in their ability to switch fuel sources, and production infrastructure is thought to be operating near capacity. Also, significant lead time is required in order to bring additional domestic or foreign natural gas supplies to market as well as expand pipeline capacity to alleviate transmission bottlenecks. Limited short-term price responsiveness means that natural gas prices will be highly sensitive to market factors such as weather swings or supply disruptions. Inelasticity is characteristic of many energy commodities. However, analyses of natural gas volatility relative to other commodities have ranked it among the highest.<sup>1</sup> Electricity has been the only commodity group with price volatilities consistently higher than those of natural gas.

In the absence of much real-time supply and demand data such as production, natural gas wellhead productive capacity, or natural gas consumption volumes, market participants look to natural gas prices as a barometer for current market conditions. Volatile prices create uncertainty and financial risk in the market and may increase the cost of capital, causing pipeline and other infrastructure investment to be more expensive.

The purpose of this paper is to address whether or not natural gas prices have been more volatile in recent years and identifies potential market factors that may contribute to price volatility. The analysis found that although there is no consistent increasing or decreasing trend in natural gas spot price volatility at the Henry Hub, there is a seasonal pattern with colder months exhibiting considerably higher volatility levels. Also, the analysis indicates that price volatility tends to vary between market locations. Furthermore, the relative level of natural gas in storage has a significant impact on price volatility. When natural gas in storage is high or low compared with the 5-year average level, price volatility at the Henry Hub increases. This effect is exacerbated during the months of the year surrounding the beginning and end of the heating season when storage levels are typically at the highest and lowest levels. Finally, this analysis shows that, even with relatively low levels of volatility, changes in the natural gas price level can have large impacts on the market.

---

<sup>1</sup> See the following articles: Energy Information Administration, *Derivatives and Risk Management in the Petroleum, Natural Gas, and Electricity Industries*, October 2002; Also, Henning, B., Sloane, M., and deLeon, M., *Natural Gas and Energy Price Volatility*, American Gas Foundation, October 2003.

## Market Factors That Affect Natural Gas Supply and Demand

Natural gas prices equilibrate market supply and demand. Significant changes in supply and demand over a short period often result in large price movements needed to bring supply and demand back into balance. So considerable natural gas price volatility reflects both a reality and perception that a significant shift in supply and demand conditions has occurred. This section will discuss some of the more common causes of significant short-term changes in supply and demand, which can result in large price movements.

Factors on the supply side that may affect prices, and hence volatility, include variation in natural gas storage, production, imports, or delivery constraints. Of these, storage levels receive a high amount of attention because of the physical hedge it provides during high demand periods. Also, working gas in storage often is viewed as a barometer of the supply and demand balance in the market. Below-average storage inventories of natural gas, for example, may create a perception of supply tightness, which places upward pressure on prices. Similarly, net changes in stocks are indicative of the relative balance between supply and current consumption during the period.<sup>2</sup> For this reason, above-normal withdrawals from storage or below-normal injections to storage, compared with the equivalent time period in past years, may cause short-term spikes in natural gas prices.

Disruptions caused by severe weather, operating mishaps, or planned maintenance can also cause short-term tightness in natural gas supply. In the summer of 2005, hurricanes along the U.S. Gulf Coast caused more than 800 billion cubic feet (Bcf) of natural gas production to be shut in between August 2005 and June 2006. This is equivalent to about 5 percent of U.S. production over that period and about 22 percent of yearly natural gas production in the Federal Gulf of Mexico. As a result of these disruptions, natural gas spot prices at times exceeded \$15 per million Btu (MMBtu) in many spot market locations and fluctuated significantly over the subsequent months, reflecting the uncertainty over supplies.

On the demand side, temperature changes tend to be one of the strongest short-term influences. During cold months, residential and commercial end users consume

---

<sup>2</sup> "Current consumption" refers to natural gas burned for end-use demand and for lease and plant fuel use. It excludes natural gas flows into or out of storage.

more natural gas for heating needs, which places upward pressure on prices. If unexpected or severe weather occurs, the effect on prices intensifies because supply is often unable to react quickly to the short-term demand response, especially if the natural gas transportation system is operating at full capacity. Under these conditions, prices must rise high enough to reduce the demand for natural gas. Temperatures also have an effect on prices in the cooling season as many electric power-generating plants used to produce incremental supplies to meet air conditioning needs are fueled by natural gas. Therefore, hotter-than-normal temperatures during the summer can lead to more natural gas supplies feeding natural-gas-fired power generation. This effect may reduce natural gas available for storage and increase price pressure during the winter months when inventories are relied upon to meet heating demand.

The prices and market conditions for related fuels also have an effect on natural gas markets. In the United States, most baseload electricity generation is delivered from coal, nuclear, and hydroelectric power stations. Because natural gas tends to be a higher-cost fuel, natural-gas-fired power stations more typically are used to cover incremental power requirements that arise during times of peak demand or during sudden outages of baseload capacity. However, an increase in price or a disruption in supply in any one of the competing fuel markets can spark an increase in natural gas demand. For example, hydroelectric generation went through a relatively steep decline in the late 1990s owing to droughts in the West. The supply disruption led to a 40-percent decline in hydroelectric generation between 1997 and 2001. During the same period, natural-gas-fired generation increased 33 percent as there was spare capacity and these facilities were more flexible and better positioned than coal-fired plants to respond to the deficit in electricity supply. Additionally, natural gas competes with other fuels for industrial customers. There are dual-fired facilities such as manufacturing and electric generation plants that can switch between residual or distillate fuel oil and natural gas. Even where consuming units are single-fuel fired, fuel switching may be achieved on an aggregate basis as the industrial load or electric generation shifts from units of one fuel to another. Hence, as the prices of petroleum products increase, the industrial sector increases natural gas demand.

In general, the United States has, over the past decade, seen an increase in demand for natural gas for power generation as more natural-gas-fired power plants were built for load control and environmental reasons. According to Energy Information Administration (EIA) data, natural gas deliveries to electric power consumers in the United States increased 526 Bcf or nearly 10 percent between 2001 and 2005, which is the largest increase in

natural gas use by any consumption sector during this time period. Additional natural-gas-fired equipment increases the demand in an area, placing upward pressure on prices, but also allows the supplier to more easily meet consumer electric generation needs on peak demand days.

Lastly, economic activity is a major factor influencing natural gas markets. When the economy improves, the increased demand for goods and services from the commercial and industrial sectors generates an increase in natural gas demand. This is particularly prevalent in the industrial sector, which is the leading consumer of natural gas as both a plant fuel and as a feedstock for many products such as fertilizer and pharmaceuticals. Additionally, natural gas is consumed by oil refineries and methanol plants so increases in demand for refined products increases natural gas use. This consumption increases over time as the number of new vehicles on the road and miles driven increases. Because certain industrial customers are so dependent on natural gas, high and volatile natural gas prices may have a detrimental effect on plant economics as costs become prohibitively expensive or unpredictable. Further, economic growth, which increases personal disposable income, can also lead to an increase in residential demand.

## Scope and Methodology

This paper examines natural gas spot prices mainly between January 1994 and December 2006. The “spot price” represents the price for natural gas sales contracted for next day or weekend delivery and transfer at a given trading location. The Henry Hub is the primary trading location used to examine volatility in this paper because it is a centralized point for natural gas trading in the United States and is often a representative measure for wellhead prices. To show geographic differences in volatility, spot prices in New York City and in Chicago were analyzed.<sup>3</sup>

Historical price volatility is the primary measure used in this paper. It is defined as the standard deviation of daily relative changes in price. A natural log transformation is used to calculate the daily relative price change,  $\Delta p_t$ , for trading day  $t$  (Equation 1).

$$\Delta p_t = \ln(p_t / p_{t-1}) \tag{Equation 1}$$

<sup>3</sup> Source: *NGI's Daily Gas Price Index*, Intelligence Press. The spot price at Transco Zone 6 for New York delivery is used as the New York City spot price. The time frame of analysis for New York City only covers March 1998 through December 2006 owing to the availability of data from Natural Gas Intelligence, Inc. The Chicago price is the reported Chicago citygate price.

Volatility is calculated by multiplying the standard deviation of the daily logarithmic price changes,  $\Delta p$ , for all trading days within a certain time period by the square root of the number of trading days within the time period,  $N_T$  (Equation 2).<sup>4</sup>

$$Volatility_T = \sqrt{\frac{\sum_{t=1}^{N_T} (\Delta p_t - \Delta \bar{p})^2}{N_T - 1}} * \sqrt{N_T} \quad (\text{Equation 2})$$

This paper examines volatility on an annual, monthly, and weekly basis using daily settlement prices. The price changes are calculated for every trading day within the period so the number of trading days,  $N$ , is 252 for annual volatility, 21 for monthly volatility, and 5 for weekly volatility.

This analysis investigates various factors that may influence the level of natural gas price volatility. In order to establish these relationships quantitatively, a regression analysis was performed that attempted to determine the effect of natural gas storage levels, seasons, prices, and heating degree days on weekly gas price volatility. The variables included in this regression were limited to data that are collected and reported on a weekly basis. Consequently, other potential determinants, such as economic growth and productive capacity, were not included in this portion of the analysis because they are not collected on a weekly basis.

The initial model specification to estimate *weekly* spot price volatility at the Henry Hub establishes weekly volatility for any given week,  $t$ , as a function of the Henry Hub spot price level, lagged Henry Hub spot price level, relative storage level, heating degree-days, and calendar month (Equation 3).<sup>5</sup> A natural log transformation of the variables allows us to examine the relationships in terms of elasticity. For example, the coefficient,  $\beta_1$ , is the elasticity of weekly volatility with respect to the average Henry Hub spot price, and  $\beta_2$  is the elasticity of weekly volatility with respect to the average Henry Hub spot price in the previous week.

<sup>4</sup> Previous studies of volatility have found that this specification may overstate the actual volatility over a particular time period because of the impact of non-trading days.

<sup>5</sup> Sources: Storage Level: Energy Information Administration, *Weekly Natural Gas Storage Report*; Heating Degree-Days: National Oceanic and Atmospheric Administration (NOAA), National Weather Service Climate Prediction Center. Heating degree-days (HDD) are an approximate measure of temperature and are used in this analysis to reflect demand for energy used to heat buildings. The HDD used in this paper are gas-weighted to account for the number of residential customers using natural gas (Census 2000 data).

$$\begin{aligned} \ln(Volatility_t) = & \beta_0 + \beta_1 \ln(HHspot_t) \\ & + \beta_2 \ln(HHspot_{t-1}) \\ & + \beta_3 AboveAveStorage_t \\ & + \beta_4 \ln(StorageDifference_t) \\ & + \beta_5 Heating * \ln(HDDRatio_t) \\ & + \beta_6 Feb \\ & + \dots + \beta_{16} Dec + e_t \end{aligned} \quad (\text{Equation 3})$$

where:

- $Volatility_t$  = Weekly price volatility of the Henry Hub spot price as defined in Equation 2.
- $HHspot_t$  = Average weekly Henry Hub spot price.
- $AboveAveStorage_t$  = Dummy variable equal to 1 if the weekly storage level is above the 5-year average.
- $StorageDifference_t$  = The absolute value of the percentage difference between the weekly storage level and the 5-year average.<sup>6</sup>
- $Heating$  = Dummy variable equal to 1 if the week falls within the heating season (November through March).
- $HDDRatio_t$  = The number of weekly gas-weighted heating degree-days divided by the normal number of heating degree-days for the equivalent week.<sup>7</sup>
- $Feb$  through  $Dec$  = Dummy variables representing month of the year.
- $e_t$  = the error term for week  $t$ .
- $\beta_i$  = Coefficients to be estimated ( $i = 0, 1, \dots, 16$ ).

Augmented Dickey-Fuller (ADF) tests were used to analyze the stationarity of the model, a requirement for time series regression. Because of the nature of time series data, it is logical that the effects from a random shock to an economic variable may carry over from one time period to the next. Stationarity asserts that the effects of a random shock dissipate over time. In contrast, a nonstationary time series has lasting effects resulting from a random shock in the time series. A key implication of nonstationarity is that analysis of nonstationary time series may produce spurious results. In other words, the  $R^2$  may be high with significant  $t$ -

<sup>6</sup> The absolute value allows for the use of a natural log transformation.

<sup>7</sup> The normal number of heating degree-days is calculated by NOAA using the period, 1971-2000.

statistics for the coefficients, but the results are meaningless. The ADF test checks for the existence of a unit root in the data series, which indicates nonstationarity. The ADF statistics, presented in Appendix A, show that the natural log of weekly volatility, the natural log of the *StorageDifference* variable, and the natural log of the heating degree-day variable are stationary time series. The natural log of the Henry Hub spot price, however, fails to reject the unit root hypothesis. This issue is corrected by representing the Henry Hub price as the ratio of the Henry Hub price to the lagged price, which de-trends the variables and proves to be stationary.<sup>8</sup>

The model can be rewritten to account for the modified form of the *HHspot* variable (Equation 4). The model is estimated using a first order autoregressive error model, AR(1). The AR(1) model is useful for estimating linear regression models with time series data because it corrects for autocorrelation. With autocorrelation, error terms are not independent, which may lead to inefficient estimators and invalid standard errors and test statistics. In the AR(1) error model, the error term,  $e_t$ , becomes  $\rho e_{t-1} + \varepsilon$ .

$$\begin{aligned} \ln(\text{Volatility}_t) = & \beta_0 + \beta_1 \ln\left(\frac{HHspot_t}{HHspot_{t-1}}\right) \\ & + \beta_2 \ln\left(\frac{HHspot_{t-1}}{HHspot_{t-2}}\right) \\ & + \beta_3 \text{AboveAveStorage}_t \\ & + \beta_4 \ln(\text{StorageDifference}_t) \\ & + \beta_5 \text{Heating} * \ln(\text{HDDRatio}_t) \\ & + \beta_6 \text{Feb} \\ & + \dots + \beta_{16} \text{Dec} + e_t \end{aligned} \quad (\text{Equation 4})$$

The model was estimated with data from January 1999 to December 2006. Since weekly storage data are not available prior to 1994, it is impossible to calculate a previous 5-year storage average for data earlier than January 1999, which is necessary for the *StorageDifference* variable. Results of the model are presented in Appendix B and will be referred to within the text of this paper. These estimates might be biased since it is reasonable to believe that volatility, spot prices, and storage inventories are jointly determined. However, the results of the model are a good starting point in the

examination of volatility, and lead to interesting topics for future research.

As noted in the previous section, examining volatility levels alone to address price risk may not provide a complete picture of the impacts that expanding daily price movements have on market participants. In order to address this issue, the absolute changes in daily price were examined by using the mean absolute deviation (MAD) for a given time period. For the purposes of this paper, MAD is defined as the mean of the absolute value of changes in daily settlement prices over a given period. Since this measure looks at changes in price magnitude, it is reported in dollar amounts and is useful to contrast with volatility, which is based on percent changes.

## Results

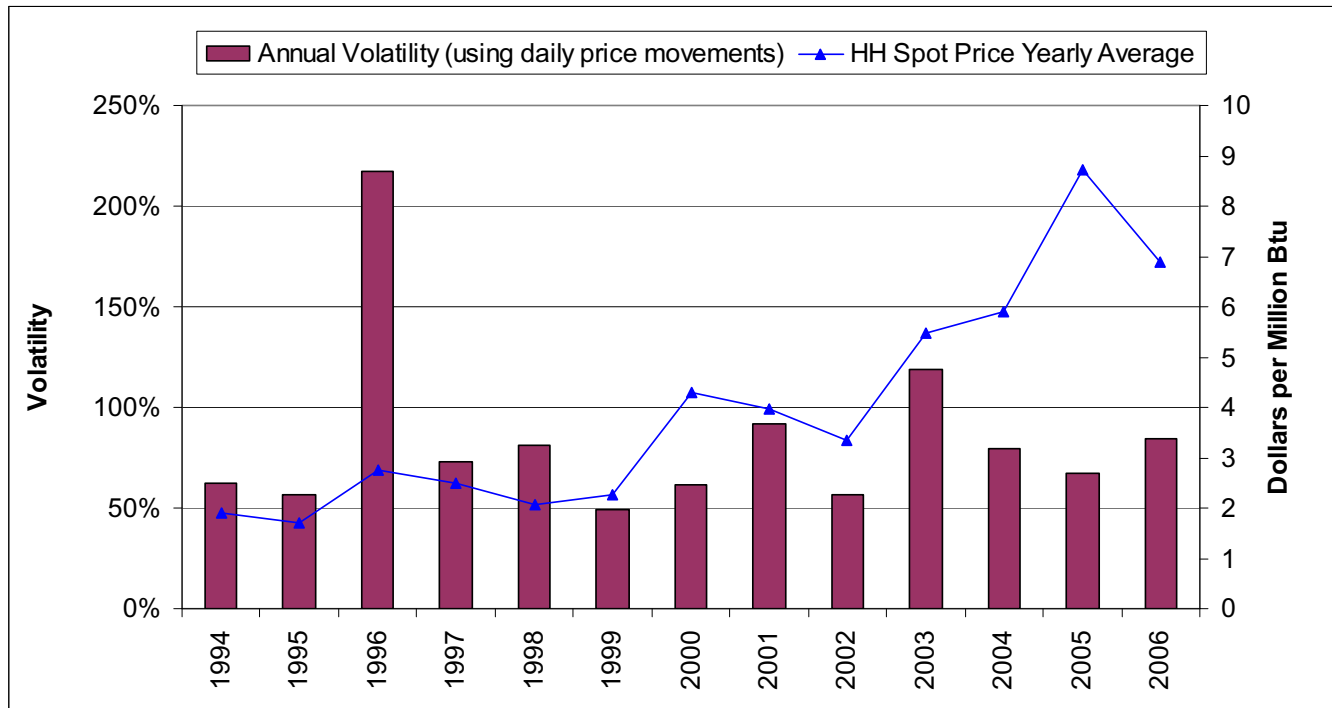
*Result 1: Annual price volatility at the Henry Hub has been high for the past decade, but it does not exhibit a consistent increasing or decreasing trend.*

An examination of daily settlement prices at the Henry Hub shows that annual volatility has fluctuated between 49 percent and 218 percent since 1994.<sup>9</sup> Although the individual values each year are high relative to some commodities, the level of annual volatility does not seem to exhibit a clear trend in recent years (Figure 2). There is clearly an overall increasing trend in the Henry Hub average price, yet the upward price movements are not reflected in the annual volatility levels. This is not surprising because, as discussed earlier, price levels and price volatility are two distinct concepts. However, a constant volatility at higher prices results in a greater dollar value price change at those higher prices.

<sup>8</sup> For additional information on unit root tests and nonstationary processes, see Green, William H., *Econometric Analysis, Third Edition*. New Jersey: Prentice Hall Inc, 1997. Chapter 18.3.

<sup>9</sup> In 1996, the historical annual volatility reached 218 percent partly because of a brief period in early February that included a price spike to more than \$14 per MMBtu. Excluding data for the first two weeks of February 1996 would yield an estimate of 109 percent for the year.

**Figure 2: Annual Volatility and the Henry Hub Yearly Spot Price Average**



*Result 2: Monthly price volatility at the Henry Hub exhibits seasonality.*

When the daily price volatility data are examined on a monthly basis, several trends begin to emerge. For the time period analyzed in this paper (January 1994 to December 2006), there is a degree of seasonality in the monthly volatilities (Figure 3).<sup>10</sup>

The highest volatility levels in any year tend to occur from October through February. December through February are generally the coldest months of the year, and demand for natural gas rises sharply during these months as heating needs dominate the market. As discussed earlier, the inelasticity of natural gas supply and demand can cause large price swings in response to market factors such as cold weather. Although November is the warmest winter month, it is the first month of the heating season, and fluctuations in demand or supply are not necessarily met readily with storage volumes because the natural gas in storage has to serve the entire winter season, whose

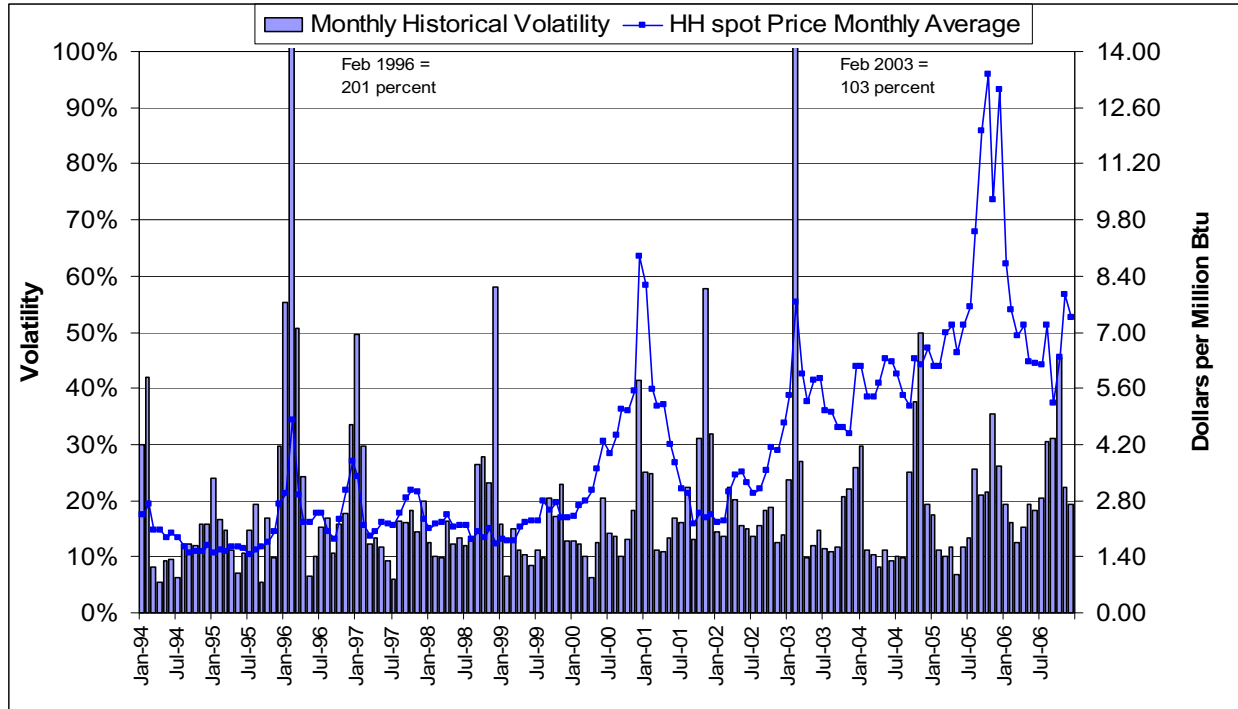
overall severity is unknown this early in the winter season. So the prices may be more erratic than in a time period with similar temperature-related demand. The high volatility levels in October are likely because October is the last month of the refill season. There may be increased competition from storage facilities looking to meet end-of-season refill goals as well as increased anticipation regarding the upcoming heating season.

Calculating the average volatilities by month further demonstrates that price volatility follows temperature patterns with the coldest months exhibiting higher volatilities (Table 1). The monthly volatilities averaged by month for the winter heating season months, except for March, ranged from 25 to 27 percent.<sup>11</sup> The shoulder months of October and March, which border these 4 cold months, were next highest. The more mild spring and summer months exhibited the lowest average levels of price volatility. The winter months also exhibited the largest coefficients of variation suggesting that these averages are influenced more by outliers compared with the warmer months.

<sup>10</sup> Similar to the pattern in the annual volatility estimates, in February 1996, the monthly volatility is 201 percent largely owing to a price spike in the beginning of the month. If the first two weeks of February are eliminated from the calculation, monthly historical volatility in this month decreases to 60 percent.

<sup>11</sup> The heating season for natural gas is November through March.

**Figure 3: Monthly Volatility and the Henry Hub Monthly Spot Price Average**



Note: The maximum volatility level shown on the graph is 100 percent, which truncates the February 1996 observation of 201 percent and the February 2003 value of 103 percent. This was done to allow more detail to be seen in the other portions of the volatility series.

Wilcoxon rank-sum and median two-sample tests were used to determine whether the differences in volatility across months were statistically significant. The data were separated into three categories based on the average monthly volatilities in Table 1: (1) high volatility months (January, February, November and December), (2) shoulder months (March and October) and (3) low volatility months (April through September). The comparisons between the high and low categories determined that volatility in the high months is significantly greater than that of the low months at a 5-percent level of significance for both the Wilcoxon rank-sum and median two-sample tests. However, comparisons involving the shoulder months were not conclusive. Neither the Wilcoxon rank-sum test nor the median two-sample tests were able to determine that the shoulder months were significantly different from the high months at a 5-percent level of significance. Comparing the shoulder month category and the low category found that volatility in the shoulder months is significantly different than that of the low category, according to the Wilcoxon rank-sum test, but not according to the median two-sample test at a 5-percent level of significance.

**Table 1: Average Monthly Volatility by Month**

Month	Average Monthly Volatility	Coefficient of Variation
January	25.38%	52.98%
February*	27.45%	99.06%
March	16.49%	70.37%
April	12.63%	42.30%
May	11.38%	31.47%
June	12.87%	30.06%
July	12.67%	30.72%
August	16.32%	35.42%
September	17.05%	43.89%
October	22.79%	43.76%
November	24.75%	58.43%
December	26.74%	47.29%

Note: The darkest shading represents the months with the highest volatility levels. The lighter shading indicates the two shoulder months.  
 \*The calculated February values exclude the first two weeks in February 1996. If these weeks were included average February volatility would be 38 percent with a coefficient of variation of 69 percent.

These results, which are presented in Appendix C, provide evidence that volatility levels in the high months differs from the low months to a statistically significant

degree. However, the conflicting results of tests on the shoulder months likely reflect the presence of outliers in the data during the shoulder months.

The results of the time-series analysis of weekly volatility data corroborate a pattern of seasonal variation in volatility levels, although the results suggest that the relationship is more closely associated with storage dynamics than with temperature patterns. The coefficient on the variable representing heating degree-days during the heating season, a proxy for temperature, was found to be insignificant. Additionally, the AR(1) estimation of Equation 4 shows that weekly volatility levels during the months surrounding the start and the end of the heating season, when storage levels are at the highest and lowest levels, respectively, are statistically different than volatility during the month of the omitted dummy variable, January.<sup>12</sup> None of the dummy variables for other months, except for July, were significant at the 0.05 level.

The coefficients on the March and April variables are negative and indicate that weekly volatility in these months is 30 and 34 percent lower than in January, respectively.<sup>13</sup> During these months, the peak winter demand is generally complete, and there is less uncertainty regarding whether or not there will be sufficient supplies for heating needs. Although winter-like temperatures sometimes persist into April or summer-like temperatures arrive early, it is during this month that attention tends to switch from storage withdrawals towards storage injection as shown by average working gas volumes in each month (Table 2). Furthermore, there is generally a low demand for air conditioning needs and natural-gas-fired electric generation during the spring months, allowing for natural gas storage to build.

The coefficients on the October and November variables are positive and indicate that weekly volatility in these months is 46 and 47 percent higher than in January, respectively. In contrast to March and April, October and November are the months in which working natural gas inventories are typically highest. During October, participants attempt to achieve their storage targets for the upcoming heating season. Storage capacity owners may be competing heavily to inject natural gas before the winter season and have diminished flexibility in meeting their goals because time is running out. Competition from

lingering cooling load needs or early-season heating demand is also common. This additional need for natural gas adds to tightness in market supply, contributing to higher volatility.

**Table 2: Average Volume of Working Natural Gas in Storage at the End of the Month, 1999 – 2006**

Month	Average Volume of Working Gas in Storage by Month (Million Cubic Feet)
January	1,798,139
February	1,344,223
March	1,138,762
April	1,290,824
May	1,638,813
June	1,995,669
July	2,295,990
August	2,565,378
September	2,883,008
October	3,084,828
November	2,975,213
December	2,486,229

Note: Shaded lines represent months for which the monthly dummy variable was statistically significant in the AR(1) error model.

As the heating season begins in November, storage activities turn toward withdrawals. Decisions made during this month impact the volumes remaining in storage for the rest of the heating season. Owners of natural gas in storage often are reluctant to withdraw significant amounts of natural gas in November because of the uncertainty regarding the availability of supplies later in the winter. Natural gas in storage helps suppliers ensure their ability to meet contractual or regulatory requirements in the following months.<sup>14</sup>

These results are generally consistent with previous analysis of price volatility in natural gas markets, however they are not strictly comparable as this model tests for differences from January volatility. The results from the AR(1) model also may seem to contradict the calculations of average monthly volatilities that showed January having a greater level compared with November and December. However, the AR(1) model excludes years prior to 1999. As it turns out, January volatility

<sup>12</sup> If all the dummy variables are included in the model, they would fully explain the variation in volatility and the coefficients would sum to one. A common solution to this issue is to omit one dummy variable. Results of the AR(1) error model are available in Appendix B.

<sup>13</sup> Because of the log transformation, the coefficient is not the percent change in weekly volatility for each month compared to January. To

calculate this you must use:  $(e^{\beta_x} - 1) \times 100$

<sup>14</sup> The use of natural gas from storage is not entirely certain because it typically requires withdrawal and transportation services, either of which may be subject to failure to perform properly, thus jeopardizing the flow of supplies when needed. However, such failure to perform is not common.

levels were relatively high during the years between 1994 and 1997, which were included in the average monthly volatility calculations.<sup>15</sup> Also, some seasonal impacts may be captured by systematic variation in the storage variable of the AR(1) model because storage levels, which exhibit a seasonal pattern, have an impact on price, as discussed in Result 4.

Overall, the results support the finding of a seasonal trend in the data, continuing to suggest that months with higher levels of market tightness and uncertainty have higher volatility levels. This finding is reasonable given that supplies may not keep pace with the increased demand during winter months, or if demand does not materialize because of mild weather. The tightening or loosening of the market during the winter has a more pronounced effect on natural gas prices because both supply and demand are relatively inelastic during this time of year. Therefore, during the winter months, natural gas prices tend to swing more in order to balance supply and demand.

*Result 3: Natural gas price volatility shows a positive correlation with the relative change in the contemporaneous spot price, yet volatility does not seem to be following a clear increasing trend similar to the spot price.*

The AR(1) model results show a relationship between weekly volatility and the relative change in the average weekly Henry Hub spot price while controlling for calendar month and storage level. The coefficient on the natural log transformation of the relative change in average weekly Henry Hub spot price is 1.16. This model indicates that a 1-percent change in the average weekly Henry Hub spot price yields a 1.16-percent change in the weekly volatility level in the same direction.

The model specifies the relation between the volatility level and movements in the average weekly Henry Hub spot price, not the price level. As explained in a previous section, the price level was not used because the time series of the Henry Hub spot price was nonstationary. Although there is a conceptual similarity between volatility and the relative change in average weekly prices as they are both measures of price change, there are certain fundamental differences between these measures used in this analysis. First, the measures reflect different aspects of price variation. The weekly volatility measures the spread or dispersion of daily price changes around the average, while the price variable is the ratio of average prices between weeks. These differences may impart to each measure dissimilar characteristics. For example, if

daily prices fluctuate around a fairly stable mean level, volatility might increase while the relative change between weeks is close to zero. Second, volatility is independent of the direction of the price change, while the ratio of average weekly prices varies directly with the price movement. Thus, both measures would increase with an increase in prices between periods. However, as prices fall, the price ratio will decline but volatility may increase, and this difference is magnified as the price movement is larger.

Graphical analyses were used to examine whether monthly price volatility is increasing, similar to the widely accepted upward trend in price level. In order to control for the effect of seasonality, the data were analyzed by calendar month to separate out the impacts that storage and temperature have on volatility (Appendix D). Slight visual trends emerge when we look at the average monthly volatilities separated by month. However, virtually no long-term trends are apparent, based on estimation of a simple trend model in which the volatility in time  $t$  is regressed on year. The relatively small  $R^2$  values estimated for each time series indicate that volatility levels in any calendar month do not exhibit stable incremental growth between years. Additionally, most of the time series for each month exhibited several large swings in direction during the time period. The results for September and October include slightly better results, such as higher  $R^2$  values. However, a conclusion of a significant time trend would require additional research beyond the scope of the current analysis.

*Result 4: As storage levels deviate more from the previous 5-year average, the level of price volatility increases.*

A graphical examination of monthly storage data suggests that when the storage inventory level moved away from the 5-year average, volatility increased albeit with a lag (Figure 4). In many years, the trend seems to be similar regardless of whether the storage level was above or below average. For example, in the winter of 2002-2003, storage levels decreased to about 40 percent below the 5-year average by February. The monthly volatility that month exceeded 100 percent. Additionally, during the heating season of 2001-2002, storage levels were 33 percent above normal on average, and volatility was 58 percent on average during that season.

The difference between storage levels and the 5-year average for a given month gives a relative measure of storage inventory. The 5-year average is a rolling average calculated from the monthly levels for those months over the previous 5 years. There are clear spikes in volatility when storage levels are relatively high or low as compared with the 5-year average. Furthermore, these

---

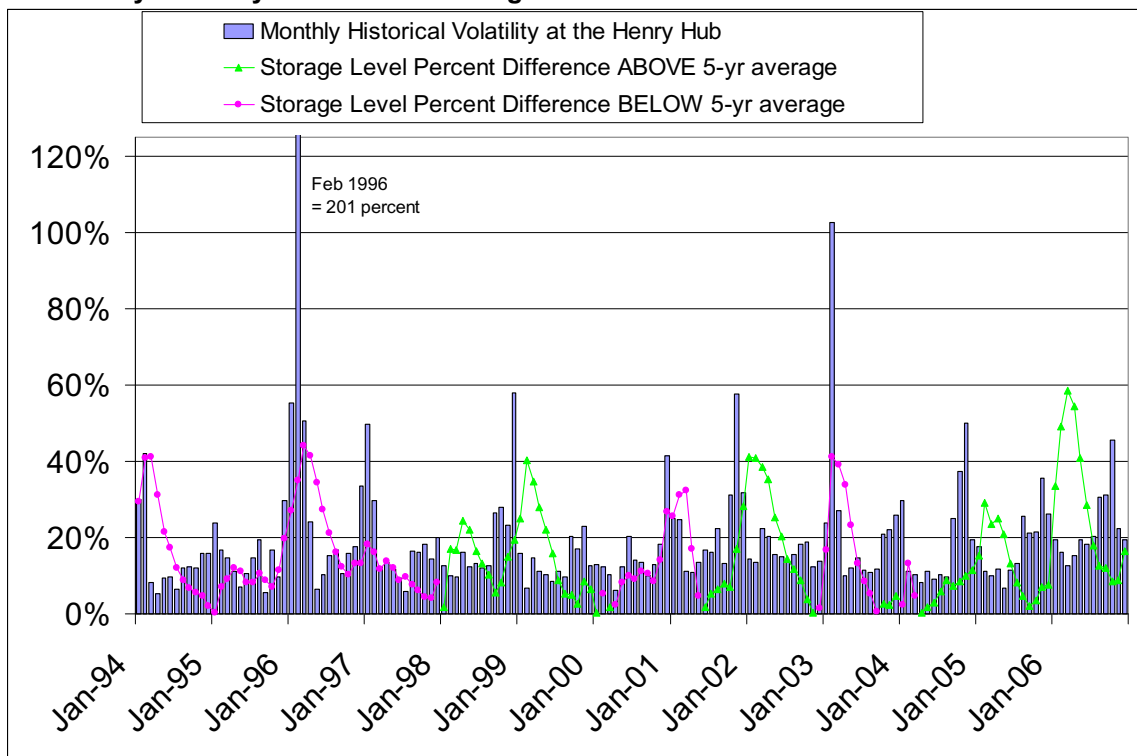
<sup>15</sup> See Appendix C for individual price volatility levels within each month of each year.

spikes in volatility seem to be larger when storage levels shift away from the 5-year average rapidly. For example, storage levels between October 1995 and February 1996 fell from 7 percent below the 5-year average to 35 percent below the 5-year average. During this time, monthly volatility increased dramatically, peaking at 201 percent in February 1996. A similar pattern occurred during the heating season of 2002-2003 when storage fell to 44 percent below the 5-year average and monthly volatility rose above 100 percent.

The monthly data suggest a relationship between relative storage levels and volatility, yet it is also clear that this relationship is rather complex. For instance, there appears

to be a slight lag between the measured volatility levels and the storage levels in many years such that the relative storage level peaks one or two months after the peak of volatility. This may suggest that volatility affects demand for storage in later time periods or that the impact of storage on volatility has an upper limit beyond which market forces take precedence. Regardless, the basic relationship sets the stage for the time-series analysis in the weekly volatility model. Areas for further research include time-lag issues, geographic differences, or whether price swings in a certain direction are associated with storage levels.

**Figure 4: Monthly Volatility and Relative Storage Levels**



The weekly AR(1) error model uses two variables to estimate the effects of storage levels on the dependent variable, weekly volatility. The *StorageDifference* variable measures the percentage difference between the weekly storage level and the 5-year average.<sup>16</sup> Since the calculation uses absolute values, a separate dummy variable, *AboveAveStorage*, is used to indicate whether the weekly storage level is above the 5-year average. The coefficient on the dummy variable is insignificant which supports our perception from the graph that changes in volatility levels are independent of whether storage is above or below the 5-year average.

The coefficient on *StorageDifference<sub>t</sub>* is 0.1206 and is significant at the 0.05 level, which indicates that, other things being equal, a 1-percent increase in the relative difference in storage compared with the 5-year average yields a 0.12-percent increase in volatility. As levels move away from the 5-year average, above or below, volatility increases slightly. This makes sense considering that working gas stocks are not only a measure of the relative balance of natural gas supplies, but storage operators add a component of competition to the market. Storage operators withdraw natural gas from storage when current supplies are insufficient to meet demand, but they also have various incentives to keep natural gas storage levels high, such as ensuring supplies for upcoming periods of high demand or economic incentives related to the future price of natural gas.

A low inventory of working gas relative to the 5-year average is consistent with tightness in the natural gas market. When storage supplies are below average in the winter (November to March), current supplies (domestic production and imports) likely are relatively low, which has led market participants to rely more heavily on storage supplies. As the storage level moves further below the 5-year average during the refill season (April to October), there is greater uncertainty regarding whether storage supplies will be sufficient to meet peak demand needs over the next year, adding upward pressure on market prices. On the other hand, working gas stocks above the 5-year average indicate that the market is relying more heavily on current supplies to either add to working gas stocks or avoid drawing stocks down. Although it often is interpreted that storage levels above the 5-year average indicate an ease in market tightness, the stocks themselves do not impact market conditions directly. It is the flows into or out of storage that have a direct impact on the market by altering the current supply or demand conditions.

---

<sup>16</sup> The *StorageDifference* variable is calculated as the natural logarithm of the absolute value of the percent difference of weekly working natural gas storage from the previous 5-year average.

As indicated by the storage variables in the time-series model as well as the month dummy variables, storage dynamics play a decidedly important role in market prices. Although this analysis provides a statistical finding of this relationship, it would be interesting to further research the seasonal differences in operator behavior.

*Result 5: Price volatility trends in various markets are mixed.*

There is no single domestic spot market in the United States, but rather a network of related spot markets in various regions. Each market has unique characteristics such as weather, proximity to supply, pipeline capacity, composition of demand, and volume of trades. Therefore, not all regional spot markets may exhibit the same trends. To this point, analysis has focused on the Henry Hub market. This section examines data for two other markets that are commonly used as benchmarks in their respective regions: New York City and Chicago.

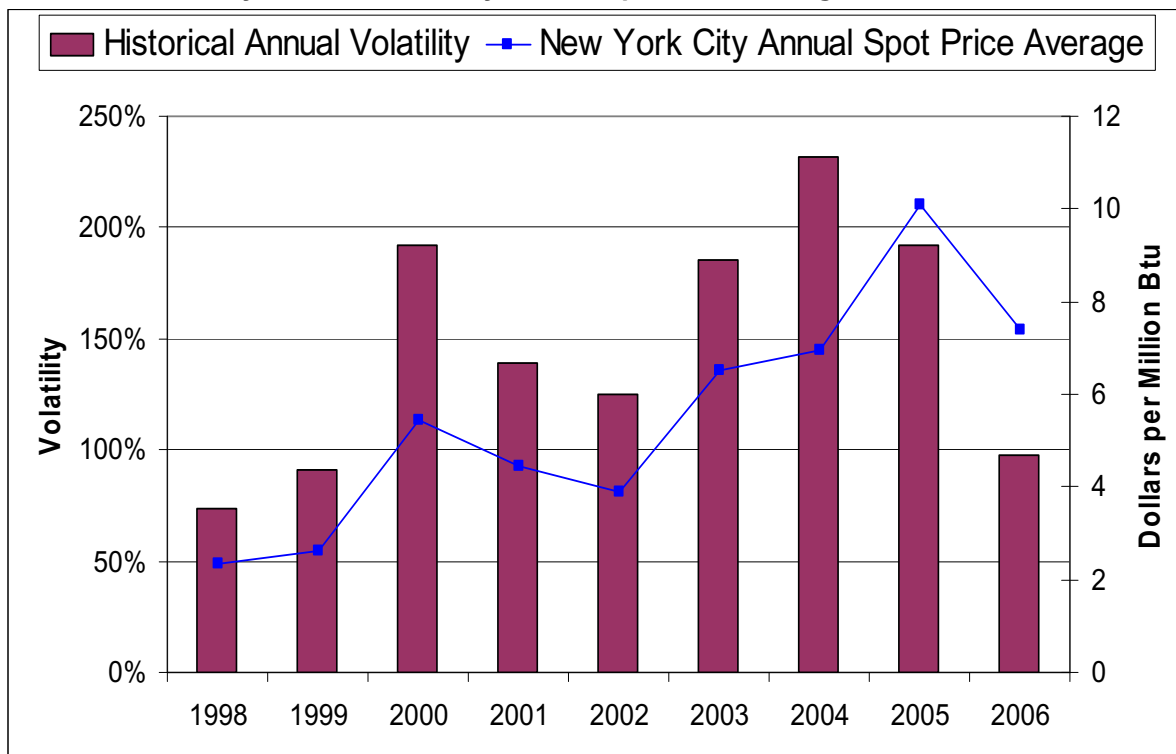
At New York City, annual volatility is significantly higher than at the Henry Hub and has been higher in recent years (Figure 5). Volatility levels at the Henry Hub rarely climbed over 100 percent during the time period analyzed, whereas volatility levels at New York City have been consistently above 100 percent since 2000 and around 200 percent between 2003 and 2005.<sup>17</sup> Another factor differentiating price patterns at this location from the Henry Hub is that, generally, the level of volatility for a given year seems to follow the same movement as the average spot price in that year. The only anomaly in this trend was in the year 2005 when volatility fell from 231 percent to 191 percent, but the average spot price increased from \$6.96 to \$10.08 per MMBtu.

Several factors differentiate New York City from locations such as the Henry Hub and may explain the relatively exaggerated volatility levels. First, New York City is separated geographically from natural gas fields and production; so much of the natural gas coming to market has to be transported over longer distances. This makes the price vulnerable to congestion or disruptions in the supply chain.

---

<sup>17</sup> In 2004, the annual historical volatility at New York City is 232 percent partly because of two trading days in January 2004 where the spot price rose to \$44.81 and \$27.96 per MMBtu. Excluding these two days reduces the volatility level to 173 percent in 2004. The yearly spot average in 2004 would decrease from \$6.96 per MMBtu to \$6.72 per MMBtu.

**Figure 5: Annual Volatility and New York City Annual Spot Price Average**



Localized capacity constraints have been an issue in the New York City area as the existing pipelines operate at full capacity during peak periods.<sup>18</sup> Demand is hugely influenced by the dense population and cold winters, and the market becomes tighter when the area experiences extreme temperatures. Lastly, most power generation in New York City is fired by natural gas or residual fuel. The need for more power generation capacity and the attractiveness of natural-gas-fired plants for environmental and economic reasons has created an increased reliance on natural gas. The expanded use of natural gas for power generation adds to market pressures during times of peak electricity demand.

At Chicago, the pattern is similar to the Henry Hub where annual volatility has been consistently high but does not seem to have a long-term increasing or decreasing trend (Figure 6). Like the Henry Hub, volatility levels have typically been below 100 percent and they are difficult to predict from year to year based on past observations.<sup>19</sup>

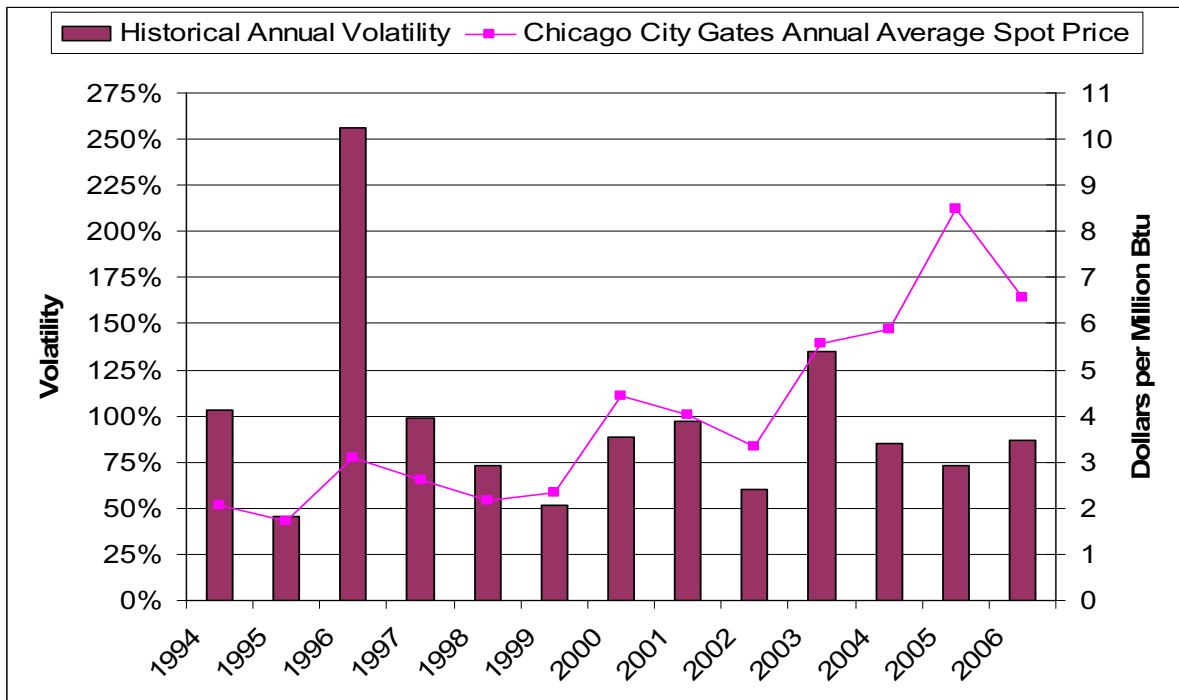
The lower price volatility in the Chicago market reflects the impact of relatively more supplies being available in the general vicinity. Chicago markets have the advantage of a major trading hub and large capacity pipeline systems in the area. The relatively plentiful supply mitigates the impact on price of shifts in demand.

The average monthly volatilities, separated by month, further show the impact of location on volatility (Table 3). The volatilities in most months are similar for Henry Hub and Chicago although volatility at the Chicago market is generally larger than at the Henry Hub in the winter months. However, volatility at New York City is often significantly higher than at the other two locations throughout the year.

<sup>18</sup> New York City Energy Policy Task Force, *New York City Energy Policy: An Electricity Resource Roadmap*, January 2004.

<sup>19</sup> In 1996, the annual historical volatility at Chicago was 256 percent partly because of two price spikes in February and March where the spot price rose to \$18.49 (over four days) and \$14.15 (over two days) per MMBtu respectively. Excluding these days reduces the volatility level to 167 percent in 1996. The yearly spot average in 1996 would decrease from \$3.10 per MMBtu to \$2.88 per MMBtu.

**Figure 6: Annual Volatility and the Chicago Citygate Annual Spot Price Average**



**Table 3: Average Monthly Volatility by Month, 1994 – 2006\***

Month	Henry Hub		New York City		Chicago	
	Average Monthly Historical Volatility	Coefficient of Variation	Average Monthly Historical Volatility	Coefficient of Variation	Average Monthly Historical Volatility	Coefficient of Variation
January	25.38%	52.98%	87.02%	65.50%	31.94%	81.03%
February	27.45%	99.06%	48.64%	83.14%	38.41%	120.52%
March	16.49%	70.37%	23.37%	54.66%	26.12%	155.78%
April	12.63%	42.30%	14.57%	34.36%	13.81%	47.77%
May	11.38%	31.47%	13.37%	36.40%	12.33%	36.05%
June	12.87%	30.06%	18.25%	26.84%	13.47%	30.50%
July	12.67%	30.72%	21.13%	67.63%	12.96%	33.99%
August	16.32%	35.42%	23.58%	71.68%	15.16%	35.76%
September	17.05%	43.89%	20.68%	33.81%	16.68%	51.19%
October	22.79%	43.76%	28.52%	34.97%	24.10%	46.72%
November	24.75%	58.43%	37.40%	51.99%	25.90%	61.87%
December	26.74%	47.29%	48.71%	51.44%	28.29%	64.96%

Note: The darkest shading represents the months with the highest volatility levels. The lighter shading indicates the two shoulder months.

\*New York City estimates only include data from 1998-2006.

*Result 6: Even under constant or relatively low levels of volatility, price risk may increase as the price of natural gas increases.*

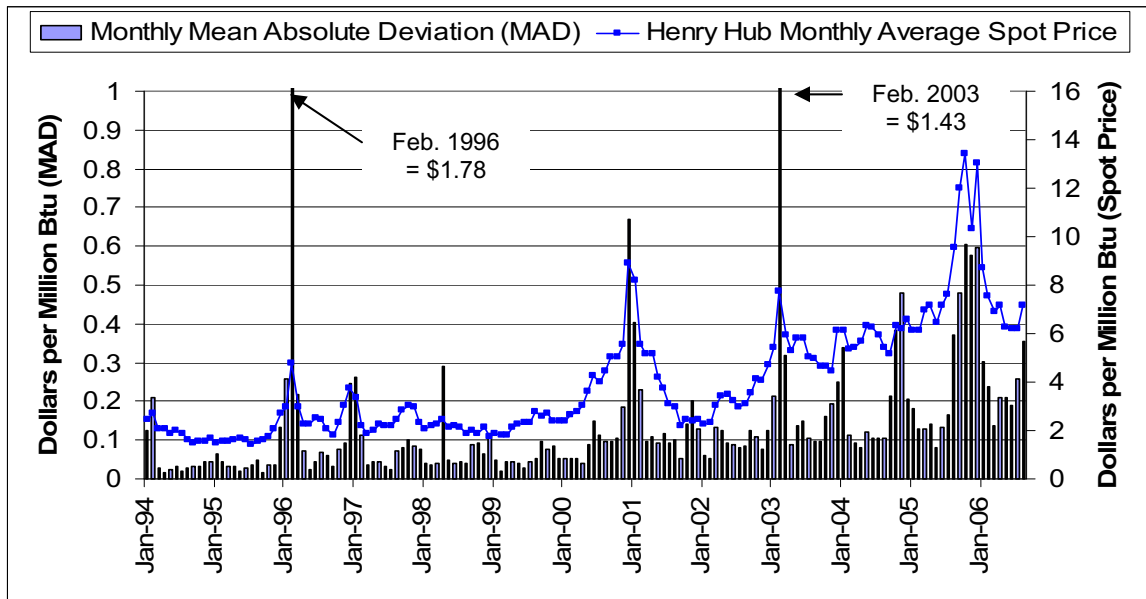
This section presents an examination of price variation in the context of absolute changes in prices as opposed to percent changes in price, or volatility. The range of

potential costs to buyers and sellers depends on the range of possible price changes. The volatility level alone does not provide a clear measure of the impact of the price changes. The same level of volatility under low prices has a smaller range of potential costs to participants in the market. As prices have increased, the growing magnitude of daily price fluctuations has increased the potential impact on customer costs.

An examination of the MAD of natural gas spot prices at the Henry Hub shows a different picture than the one based on volatilities (the analysis in this section is based on constant 2006 dollars). The MAD is calculated by taking the mean of the absolute daily price movements in a given time period. It therefore represents the average dollar amount that the price changed each day during that time period. In addition to a seasonal trend, the MAD has been increasing over time along with spot prices,

February 1 and March 1, when a sharp price spike from around \$3.50 per MMBtu in late January to more than \$18 per MMBtu on February 2 was followed by large price drops in subsequent days. In more recent years, the Henry Hub market has seen an increasing number of days when the spot price jumped a significant amount.

**Figure 7: Monthly Mean Absolute Deviation and the Henry Hub Monthly Average Spot Price**



particularly since 2000 (Figure 7). This differs from the volatility trend in the long run, which was independent of the level of prices. An increasing MAD means that spot prices during each month fluctuate more widely over time.

The finding of an increasing MAD over time is reinforced by calculating the number of days in each calendar year that the absolute deviation in price from the previous day exceeded 25 cents, 50 cents and \$1 from 1994 to 2006 (Table 4). In early years, there are relatively few days meeting any of these criteria. The 11 days in 1996 when the absolute deviation exceeded \$1 all occurred between

**Table 4: Number of Trading Days Meeting MAD Criteria**

Year	Number of Trading Days with Absolute Deviations Meeting the Following Criteria:			Total number of trading days
	≥\$0.25	≥\$0.50	≥\$1.00	
1994	12	6	0	254
1995	4	3	0	250
1996	47	27	11	250
1997	19	4	2	251
1998	12	3	2	251
1999	3	0	0	250
2000	35	14	6	249
2001	47	17	1	250
2002	15	0	0	249
2003	51	19	6	250
2004	58	25	5	249
2005	90	40	13	241
2006	117	39	2	249

Note: Monthly Absolute Deviations (MAD) are based on constant 2006 dollars.

Although the percentage changes have not increased with higher prices, a larger MAD still can have risk implications in the market because the range of cost uncertainty grows. A comparison of data for 2 months illustrates the impact of price and volatility on commodity costs (Table 5). In this example, the cost impact of price changes in a month with higher volatility (32 percent), yet a relatively low average spot price (\$2.42 per MMBtu), was compared with that in a month with lower volatility (19 percent), but a relatively high spot price (\$6.58 per MMBtu). Although the calculated volatility was lower in December 2004, the range of potential costs to the consumer or revenue for the supplier was almost double

that in December 2001, \$2.3 million compared with \$1.4 million. Variation in commodity prices has a similar, although opposite, impact on supplies.

Cost uncertainty can be a major form of risk. The stable trend in price volatility at the Henry Hub indicates that this market has not changed in a fundamental way as prices increased. This result, however, does not reflect the substantial impact that daily price changes may have on consumers or producers.

**Table 5: Comparison of the Impact of Price Level And Volatility On Industry-Wide Daily Commodity Costs**

Month	Monthly Volatility (%)	Monthly Mean Absolute Deviation (\$)	Henry Hub Average Spot Price (\$/MMBtu)	Lower Limit Price (\$/MMBtu)	Upper Limit Price (\$/MMBtu)	Industry Wide Daily Costs based on hypothetical volume (900,000 MMBtu/day)*		
						Lower Limit Cost (\$/day)	Upper Limit Cost (\$/day)	Range in Cost (\$/day)
Dec 2001	32%	0.13	2.42	1.65	3.20	<b>1,485,326</b>	<b>2,877,306</b>	<b>1,391,980</b>
Dec 2004	19%	0.21	6.58	5.30	7.87	<b>4,772,227</b>	<b>7,080,344</b>	<b>2,308,118</b>

\* 900,000 MMBtu/day is a hypothetical volume based on the average volume traded per day during December 2004. Volume data is not available for December 2001.

Source: *NGI's Daily Gas Price Index*, Energy Intelligence.

## Conclusions

A high degree of price volatility seems inherent in natural gas markets owing to the nature of the commodity, supply capacity constraints, and the sensitivity of peak day demands to temperatures. As prices rise, there is a common perception that volatility also is increasing. However, volatility and high prices are different aspects of market pricing.

Analysis based on statistical tests, a first-order autoregressive error model, and graphical analysis, demonstrates that volatility is a complex issue with several influencing factors. This analysis shows that although annual volatility during 1994 to 2006 does not exhibit a clear overall trend, there are several patterns within the data.

- At the monthly level, volatility is higher during the cold months when short-term demand peaks. This effect is exacerbated during the months surrounding the start of the heating season when upcoming high-demand needs are less certain and eases around the

end of the heating season, suggesting that storage dynamics have a dominant role in influencing volatility levels.

- Although there are seasonal trends in the volatility levels, there is no clear increasing or decreasing trend when the data are analyzed by calendar month.
- The change in weekly price volatility increases as storage levels move away from the 5-year average. This relationship is independent of whether storage levels are above or below the 5-year average.
- Some markets may experience even greater price volatility than at the Henry Hub, such as New York City where transportation constraints sometimes are binding during the winter.
- Finally, although volatility, which is defined on the basis of percent changes in market prices, may be stable, this trend obscures the absolute impact on costs or revenues as the market price grows. Volatility may not be increasing, but even under relatively low levels of volatility, financial risk can be large as daily price movements expand.

## Appendix A: Augmented Dickey Fuller test Results for variables in Equations 3 and 4

Variable	ADF Statistic	Reject or Fail to Reject that a unit root is present	Mean	Standard Deviation
$\ln(\text{Volatility})$	-8.68	Reject ( $\rho = .01$ )	-2.68	.607
$\ln(\text{HHspot}_t)$	-2.41	Fail to Reject	1.52	.476
$\ln\left(\frac{\text{HHspot}_t}{\text{HHspot}_{t-1}}\right)$	-13.83	Reject ( $\rho = .01$ )	.003	.085
$\ln(\text{StorageDifference})$	-4.53	Reject ( $\rho = .01$ )	-2.27	1.06
$\ln(\text{HDDRatio})$	-9.21	Reject ( $\rho = .01$ )	-0.194	0.415

Notes:  $\ln$  = natural log; Volatility = the average volatility level in a given week;  $\text{HHspot}_t$  = the average Henry Hub spot price in a given week; StorageDifference = The absolute value of the percentage difference between the weekly storage level and the 5-year average level; HDDRatio = The ratio of heating degree-days in a given week divided by the normal number of heating degree-days in the equivalent week.  
 ADF statistic critical value =  $\sim -3.4$  ( $\alpha = .05$ )

**Appendix B: Results of the AR(1) Error Model estimated by Equation 4;  
Dependent variable is the natural log of weekly volatility**

Variable	Estimate	Standard Error	t	Pr > t
<b>Intercept</b>	-2.3073	0.1626	-14.19	<.0001
<b><math>\ln\left(\frac{HHspot_t}{HHspot_{t-1}}\right)</math></b>	<b>1.1623</b>	<b>0.3191</b>	<b>5.52</b>	<b>&lt;.0001</b>
$\ln\left(\frac{HHspot_{t-1}}{HHspot_{t-2}}\right)$	-0.0493	0.3114	-0.16	0.8743
<i>AboveAveStorage</i>	-0.440	0.0874	-0.50	0.6153
<b><math>\ln(\text{StorageDifference})</math></b>	<b>0.1206</b>	<b>0.0380</b>	<b>3.17</b>	<b>0.0017</b>
<i>Heating*ln(HDDRatio)</i>	0.1490	0.2621	0.57	0.57000
<i>February</i>	-0.1983	0.1575	-1.26	0.2088
<b><i>March</i></b>	<b>-0.3626</b>	<b>0.1666</b>	<b>-2.18</b>	<b>0.0302</b>
<b><i>April</i></b>	<b>-0.4161</b>	<b>0.1688</b>	<b>-2.46</b>	<b>0.0142</b>
<i>May</i>	-0.2479	0.1696	-1.46	0.1447
<i>June</i>	-0.2141	0.1697	-1.26	0.2079
<b><i>July</i></b>	<b>-0.3734</b>	<b>0.1888</b>	<b>-1.98</b>	<b>0.0487</b>
<i>August</i>	-0.2320	0.1785	-1.30	0.1944
<i>September</i>	0.0570	0.1705	0.33	0.7382
<b><i>October</i></b>	<b>0.3818</b>	<b>0.1720</b>	<b>2.22</b>	<b>0.0271</b>
<b><i>November</i></b>	<b>0.3848</b>	<b>0.1679</b>	<b>2.29</b>	<b>0.0225</b>
<i>December</i>	0.2761	0.1548	1.78	0.0754

Note: Variables in **bold** are statistically significant at the .05 level.

R-squared = 0.3671, MSE = 0.24926, Durbin-Watson = 2.0590 (Pr<DW = .5832)

Definitions: ln = natural log; HHspot<sub>t</sub> = the average Henry Hub spot price in a given week; StorageDifference = The absolute value of the percentage difference between the weekly storage level and the 5-year average level; Heating = Dummy variable equals 1 if the week falls within the heating season; HDDRatio = The ratio of heating degree-days in a given week divided by the normal number of heating degree-days in the equivalent week of previous years.


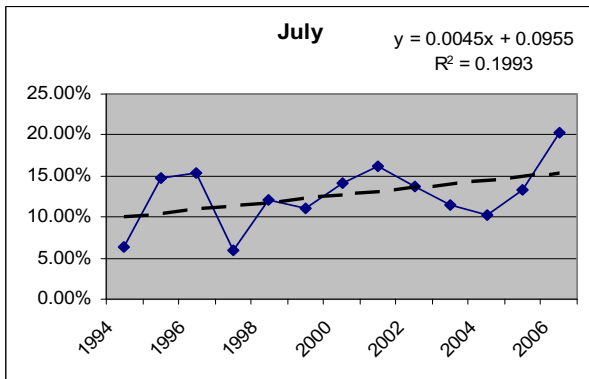

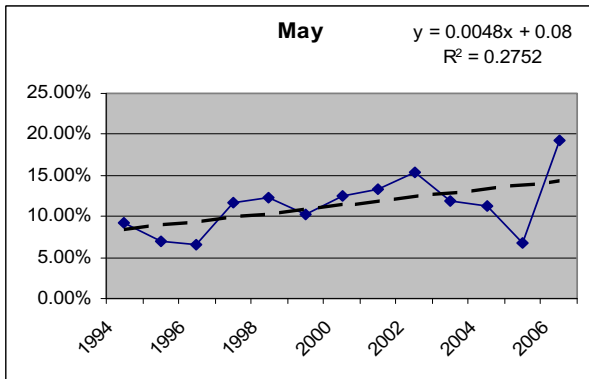
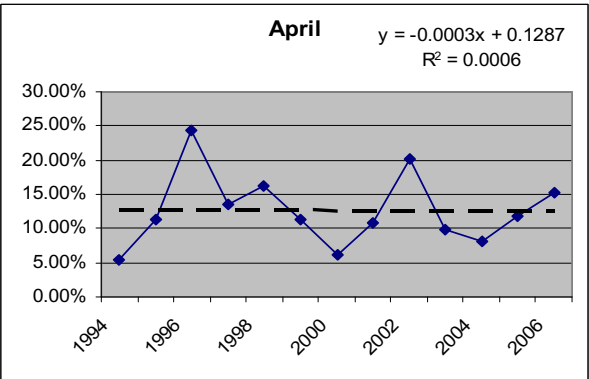
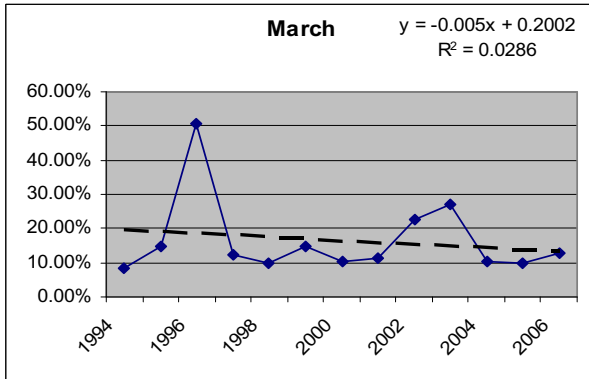

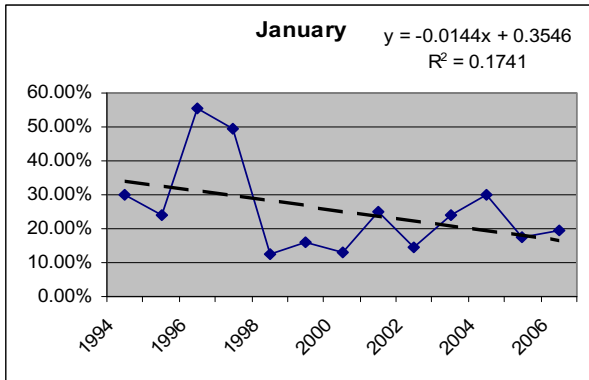
## Appendix C: Two-Sample Tests for High, Low, and Shoulder Month Categories of Monthly Volatility Levels

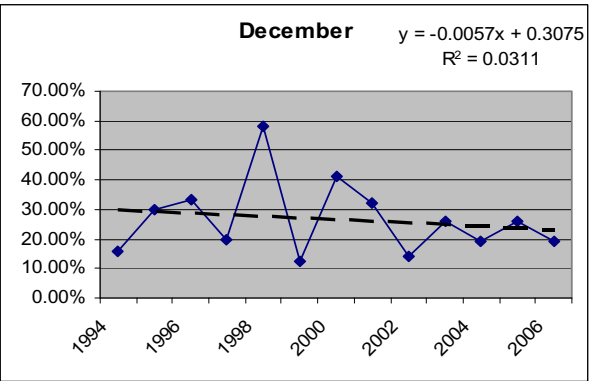
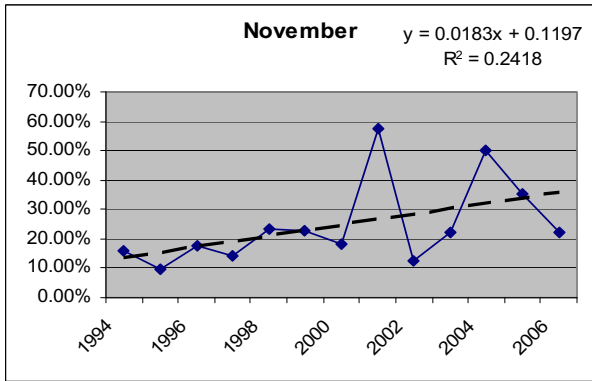

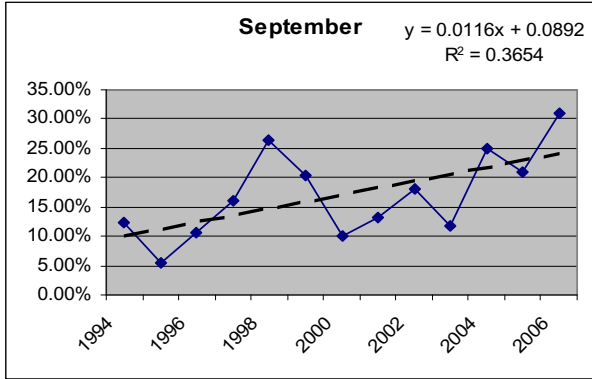
TEST	Wilcoxon Rank-Sum Test		Median Test	
	Test Statistic	Pr > t	Test Statistic	Pr > t
High/Low	4037	<.0001	36	<.0001
High/Shoulder	729	.0668	9	.1605
Shoulder/Low	1481	.0074	16	.0607

Note: The category “high” includes the months of January, February, November, and December. The category “low” includes the months from April through September. The category “shoulder” includes the months of March and October.

The Wilcoxon rank-sum test compares two independent samples to test whether they are drawn from the same population. The median test is used to test whether two samples are drawn from populations with the same median.

## Appendix D: Regressions of Average Monthly Volatilities, Separated by Month





Note: Y in the above equations represents the volatility for a given calendar month; X represents time.