CEEPR Center for Energy and Environmental Policy Research

The Weak Tie Between Natural Gas and Oil Prices

by
David J. Ramberg and John E. Parsons

10-017 November 2010

A Joint Center of the Department of Economics,
MIT Energy Initiative, and Sloan School of Management

The Weak Tie Between Natural Gas and Oil Prices

David J. Ramberg and John E. Parsons***

original draft November 2010
this draft June 2011

Abstract: Several recent studies establish that crude oil and natural gas prices are cointegrated. Yet at times in the past, and very powerfully in the last two years, many voices have noted that the two price series appear to have "decoupled". We explore the apparent contradiction between these two views. We find that recognition of the statistical fact of cointegration needs to be tempered with two additional points. First, there is an enormous amount of unexplained volatility in natural gas prices at short horizons. Hence, any simple formulaic relationship between the prices will leave a large portion of the natural gas price unexplained. Second, the cointegrating relationship does not appear to be stable through time. The prices may be tied, but the relationship can shift dramatically over time. Therefore, although the two price series may be cointegrated, the confidence intervals for both short and long time horizons are large.

* MIT Engineering Systems Division and the MIT Center for Energy and Environmental Policy Research, MIT, E19-411, 77 Massachusetts Ave., Cambridge, MA 02139 USA, E-mail: dramberg@mit.edu

** Corresponding author: MIT Sloan School of Management, MIT Center for Energy and Environmental Policy Research, and MIT Joint Program on the Science and Policy of Global Change, E19-411, 77 Massachusetts Ave., Cambridge, MA 02139 USA, E-mail: jparsons@mit.edu

This research was supported by a grant from BP Technologies Ventures Inc. as a part of the BP-MIT Advanced Conversion Research program. Helpful comments have been given by our colleagues on the MIT Future of Natural Gas study, and by colleagues at BP, Petrobras and Hess, among others.

INTRODUCTION

A number of recent academic studies have established that natural gas and crude oil prices are cointegrated.[1] These results have had an impact on analysts in the business and policy community. For example, the recent World Energy Outlook 2009, published by the International Energy Agency, reprinted a table from of one of these studies showing how an increase in the price of crude oil would be mirrored over the subsequent 12 months by a matching increase in the price of natural gas. No sooner had the results of these academic studies achieved widespread acceptance than the world witnessed a remarkable decoupling between these two prices. At the end of December 2008 the price of crude oil stood at $32.35/bbl and the price of natural gas at $5.44/mmBtu, a ratio slightly less than 6. From there the price of oil began a recovery while the price of natural gas continued to decline. At the start of September 2009 the price of crude oil stood at nearly $68.02/bbl and the price of natural gas was 1.88/mmBtu, a ratio of more than 36. At the conclusion of 2010, the price of oil reached $91.38/bbl, while the price of natural gas had recovered to only $4.23/mmBtu, yielding a ratio just above 21. So what happened to the strong tie between the prices that these studies documented? Some believe that the recent price movements reflect a permanent rupture of the old tie between the two price series—a decoupling—caused by fundamental changes in the industry. If so, then studies establishing cointegration are already outdated.

[1] The term cointegration is used to describe a certain relationship between two time series, like oil and natural gas prices, that are likely to be non-stationary—for example, because each series may grow unboundedly with the growth of the economy or with inflation, or because the variance of each series grows or shrinks with time. It is difficult to properly characterize a relationship between two non-stationary time series. Sometimes underlying the two is a single process (or combination of processes) causing them to be non-stationary. When this is the case, the relationship between the two can be represented as a line, or a linear combination. Then the two series are considered cointegrated, and the line describing their relationship is the cointegrating equation. See, for example, Hendry and Juselius (2000).

This is not the first time the natural gas price has appeared to decouple from the oil price. Throughout the 1980s and early 1990s, the United States experienced a so-called "gas bubble"–an excess supply of deliverable gas—that kept natural gas prices low relative to the then prevailing price of crude oil. The situation reversed itself in the late 1990s and early 2000s so that the price of natural gas was regularly above the level one might have predicted based on the historical relationship. Both times there was industry talk of a decoupling.

Nevertheless, throughout these periods of ups and downs, statistical analysis establishes that the two time series were cointegrated, at least until recently. How is one to rationalize these seemingly contradictory facts? What do we really mean by cointegration if the ratio of prices is shifting so consequentially across decades, sometimes in one direction, and sometimes another?

We attempt to answer these questions by elaborating on exactly what has been documented as cointegration, and putting it into context with the historically changing relationship between the two price series. We also seek to clarify what is meant when industry analysts assert that the two prices have "decoupled." These assertions are often vague and open to alternative interpretations:

- (i) the prices have *temporarily* broken away from the usual relationship to which they will later return, or,
- (ii) the prices have *permanently* broken away from the old relationship and moved into a new relationship, or,
- (iii) the two prices no longer maintain a relationship with one another at all.

Which is it? While we cannot guess the intended definition of decoupling by those who declare it has occurred, we do shed light on which of the three possible definitions of decoupling fit the data and describe the relationship between the crude oil and natural gas price series.

In this paper, we address these questions by revisiting the cointegration analyses reported by several researchers over the last twelve years. These include Serletis and Herbert (1999), Villar and Joutz (2006), Brown and Yücel (2008) and Hartley, Medlock and Rosthal (2008). Each of these papers implements a complicated set of statistical analyses of the two data series, plus a number of related conditioning variables, in order to determine if a relationship can be found with any statistical reliability, and, if so, to determine what that relationship is.[2] These analyses involve testing for a cointegrating relationship between the two variables and estimating a vector error correction model (VECM) and a conditional error correction model (conditional ECM). The results of the four papers are broadly consistent with one another, although the details of the modeling and the parameter estimates vary. In this paper we report the results of our own modeling and tests constructed along the lines of Brown and Yücel (2008). We include some more recent data than was available at the time of their analysis. We focus the discussion in the text on an exposition of the results, without walking the reader through the full set of statistical tests performed. However, these are detailed in the Appendix. Our conclusions are as follows.

[2] Villar and Joutz (2006) and Brown and Yücel (2008) directly model the relationship between natural gas and crude oil prices. Serletis and Herbert (1999) model the relationship between natural gas and fuel oil prices, among other energy prices, but do not include crude oil specifically. Hartley, Medlock and Rosthal (2008) model the relationship between natural gas and crude oil prices, but use the price of fuel oil as an intermediate step. The time windows examined vary across the studies, as does the role of exogenous conditioning variables.

Although the two price series appear to be cointegrated, this statistical fact needs to be tempered with two additional points that we think have been insufficiently emphasized in the previous literature.

First, there is an enormous amount of unexplained volatility in natural gas prices. The raw price series for natural gas—without controlling for cointegration and any explanatory variables—is approximately twice as volatile as the raw oil price series. Hence, any simple formulaic relationship between the price of oil and the price of natural gas leaves a large portion of the short-run movements in the price of natural gas unexplained. The more statistically sophisticated approach of constructing a conditional Error Correction Model, which includes the cointegrating relationship and a set of exogenous explanatory variables, and which accounts for the reversion of natural gas prices back to the cointegrating relationship, still leaves a large portion of the volatility in natural gas prices unaccounted for.

Second, the cointegrating relationship does not appear to be stable through time. Natural gas prices may be tied to oil prices, but the relationship can shift dramatically over time. While the previous literature documented that the price of natural gas seemed to be shifting up compared to the price of oil during the period 1989-2005, we show that since early 2006 this trend reversed. The period since the start of 2009 may also reflect a further decoupling of the relationship between the two series, although we may not have enough data to know yet exactly how the relationship has been redefined.

Therefore, although the two price series have been cointegrated, the confidence intervals for both short and long time horizons are large. This paper explores the nature of this apparent contradiction in an attempt to better characterize the relationship between crude oil and natural gas prices.

2. THE NATURAL GAS AND OIL PRICE RELATIONSHIP

What is the structure of the relationship of the natural gas price with the oil price? It seems natural to imagine that the price of oil and the price of natural gas would tend to rise or fall in tandem. They are both energy carriers, with one barrel (bbl) of crude oil having approximately the same energy content as six million Btu (mmBtu) of natural gas.[3] This rough logic would argue that the price of a barrel of crude oil should equal six times the price of an mmBtu of natural gas. If the price of natural gas rises by $1/mmBtu, then the price of crude oil should rise by $6/bbl.

Economists would quibble with the presumption that the ratio of prices ought to be determined exactly by the energy content equivalence. For example, Adelman and Watkins (1997) and Smith (2004) warn against valuing reserves in terms of “barrel of oil equivalent” or “gas equivalent”. The two fuels have different costs of production, transportation, processing and storage, and they serve different portfolios of end uses with only a modest overlap. The two fuels also have different environmental costs. One should expect these factors to enter into the determination of any relationship between the prices of the two commodities, and the equilibrium relationship is unlikely to match the energy content equivalence ratio. Perhaps for this reason, the industry press contains a variety of other rules-of-thumb, including the simple 10-to-1 ratio, as well as more sophisticated, burner tip parity rules. One burner tip parity rule is based on competition

[3] To be precise, 1 barrel of West Texas Intermediate crude oil contains 5.825 mmBtu.

between natural gas and residual fuel oil, while the other is based on competition between natural gas and distillate fuel oil. Both account for the transportation cost differential from the wellhead to power plants and industrial users. Both then translate the relationship back to the price of crude oil based on the typical ratio between the price of the fuel oil and the price of crude.[4] What is the formula that best describes the relationship, if any?

In fact, nothing like an energy content equivalence nor any other simply defined relationship has been persistently observed. Figure 1 shows the real spot price series in 2010 dollars for West Texas Intermediate (WTI) crude oil and Henry Hub (HH) natural gas from 1991-2010 plotted together on the same graph.[5,6] The scale for the price of natural gas is shown on the left-hand-side, while the scale for the price of crude oil is on the right-hand-side. Looking only as far back as the 1990s, the ratio of the price of oil ($/bbl) to the price of natural gas ($/mmBtu) has sometimes been as low as 2.5-to-1, and

---

[4] In Brown and Yücel (2008), the relationship generated by competition with residual fuel oil at the burner tip is given as $P_{HH,t} = -0.25 + (85\% / 6.287) P_{WTI,t}$, where $P_{HH,t}$ is the price of natural gas at the Henry Hub, and $P_{WTI,t}$ is the price of West Texas Intermediate crude oil at Cushing Oklahoma. The relationship generated by competition with distillate fuel oil at the burner tip is given as $P_{HH,t} = -0.80 + (120\% / 5.825) P_{WTI,t}$.

[5] The starting point for our data is dictated by the history of the natural gas market in the US. The Natural Gas Policy Act of 1978 gradually led to the removal of price controls on the interstate sale of natural gas in the United States. As of January 1, 1985, ceilings were removed on the sale of new gas. This was followed by the 1987 repeal of sections of the Power Plant and Industrial Fuel Use Act that restricted the use of natural gas by industrial users and electric utilities and the Natural Gas Wellhead Decontrol Act of 1989 which completed the decontrol of US natural gas prices. In addition, the Federal Energy Regulatory Commission pursued a policy of encouraging open access to natural gas pipelines, especially through Order 636. Market depth grew quickly. By April 1990, the New York Mercantile Exchange initiated trading in a natural gas futures contract.

[6] Both series are weekly day-ahead prices of commodities as sampled by Bloomberg. The natural gas prices are volume-weighted averages in $/mmBtu for delivery at Henry Hub in Louisiana. The crude oil prices are the arithmetic averages in $/bbl for West Texas Intermediate (WTI) crude oil traded at Cushing, Oklahoma. All prices were subsequently converted into real 2010 dollars.

other times as high as 36-to-1. Natural gas prices sometimes spike dramatically, without there being any noticeable change in crude oil prices.

Figure 2 shows each of the four pricing rules-of-thumb mentioned earlier. The horizontal axis is the price of oil and the vertical axis is the price of natural gas. The line for each rule gives the predicted price of natural gas as a function of the given price of crude oil. Figure 2 also shows the scatterplot of the actual combinations of crude oil and natural gas prices in our data series. Each point in the scatterplot represents a different week's pair of prices, with the week's crude oil price determining the point's location along the horizontal axis, and the week's natural gas price determining the point's location along the vertical axis. It is clear that when the oil price has been above $80/bbl all four of the rules have overestimated the natural gas price, although the 10-to-1 rule is clearly the best of the lot. In order to examine the low oil price range more clearly, Figure 3 reproduces the rule-of-thumb graphs and the scatterplot, but focused only on the lower portion of the range of oil prices, i.e., those below $30/bbl. In this range, the actual prices are arrayed widely around the residual fuel oil burner-tip-rule and the 10-to-1 rule, and only occasionally in the neighborhood of the energy-content-equivalence rule or the distillate fuel oil burner-tip-rule.

Figure 4 provides a time-series representation of the performance of each of the rules-of-thumb, graphing the prediction errors through time, i.e. the actual log natural gas price minus the predicted log natural gas price. These graphs call attention to a key problem that will undermine any simple relationship between the price of natural gas and the price of oil: there is much more volatility in the natural gas price than can be accounted for by movements in the oil price. This fact is also evident in Figure 1. The annualized volatility of the log natural gas price series is 72%, while the annualized volatility of the log crude oil price series is 39%, so natural gas was a little less than twice as volatile as crude oil.[7] Much of the volatility in the natural gas price series appears to take the form of temporary spikes in the price. These spikes have a relatively short duration.

Although no simple relationship with the oil price can account for all of the variation in the natural gas price, nevertheless, the eye can spot some rough relationship between the two price series. The price spike of 2008 is the most dramatically clear example of this, as the two price series seem to move almost in lock step. The more lasting price run-up from 2003 through 2007 also clearly reflects some tie between the two price series. Even in the time period before 2002 this rough relationship seems to show up, though with less clarity. So is the price of natural gas tied to the price of oil, or not?

Part of the problem is that a number of other variables have some short-run influence either on the price of oil or on the price of natural gas. Fluctuations in one or more of these variables can lead to the price of either natural gas or oil temporarily diverging from its long-run level. These short-run fluctuations mask whatever long-run relationship may exist, making the relationship a complicated one to properly identify. The simplest of these other variables is the seasonal fluctuation in the price of natural gas in the United States. The price of crude oil is not seasonal, so the ratio of the prices must

[7] Volatility is annualized using this formula: *Ann vol NG* = Standard deviation($log\ P_{HH,t} - log\ P_{HHt-1}$)$*\sqrt{52}$. Assuming that the time series has some element of mean reversion, then the standard deviation of the *annual* price changes is expected to be less than the *annualized* standard deviation of the *weekly* price changes. If the time series is a pure geometric Brownian motion, then annual and the annualized standard deviations are expected to be the same.

vary through the calendar year. Other variables creating short-run fluctuations are stochastic. While the WTI crude oil price is for delivery in Cushing, Oklahoma and refers to a specific type of oil produced in that region, it remains a benchmark price for crudes traded globally and it fluctuates primarily with factors affecting global demand and supply. In contrast, the price of natural gas for delivery into the Henry Hub, Louisiana is impacted much more strongly by fluctuations in supply and demand specific to the North American marketplace. These include weather events such as unexpectedly severe winter storms that cause the price of natural gas to spike, or surprisingly mild winter weather that causes the price to fall. These also include temporary interruptions to supply caused by hurricanes that shut-in production, and similar events. While the natural gas price in North America is also linked to the fluctuations in supply and demand elsewhere in the globe, prices in different regions of the world can move markedly apart from one another at times.

Identifying the underlying tie between the two prices—if any—requires filtering out the effect of these various factors. This is the challenging task to which we now turn.

3. COINTEGRATION ANALYSIS

Our analysis breaks the natural gas and oil price relationship into four components. First, there is the underlying or fundamental tie between the natural gas price and the oil price. This is called the cointegrating equation. When we say this is the fundamental tie, we mean that this is the relationship that is generally reestablished after periods in which the two prices move away from one another. Second, there is the error correction mechanism. Whenever the natural gas price has been pulled away from the fundamental tie, the price will predictably drift back towards the fundamental tie. The model estimates the rate at which this drift back occurs. Third, there are a few identifiable and recurrent exogenous factors – such as seasonality, episodic heat waves and cold waves and intermittent supply interruptions from hurricanes – that cause the natural gas price to deviate from this fundamental tie in predictable ways. The statistical analysis attempts to identify and filter out these three identifiable components. The fourth component is the residual volatility or price movement not accounted for by the first three components. These are the unexplained shocks remaining after the three identifiable components of the movements in the natural gas and oil prices have been filtered out. This residual volatility reflects the myriad temporary disruptions to the supply and demand for natural gas or oil, which, much like the identifiable and recurrent exogenous factors, pull the two prices away from the fundamental tie.

To identify these four components we implement a complicated set of statistical analyses that are described fully in the Appendix. Here we focus on just the main result, which is the estimation of this Vector Error Correction Model (VECM):

$$P_{HH,t} = \gamma + \beta\, P_{WTI,t} + \mu_t\,, \tag{1}$$

$$\Delta P_{HH,t} = a_{HH} + \alpha_{HH}\ \mu_{t-1} + \sum_{i=1}^{9} b_{HH,i}\ \Delta P_{WTI,t-i} + \sum_{i=1}^{9} c_{HH,i}\ \Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{HH,j}\ X_{j,t} + \varepsilon_{HH,t}\,, \tag{2}$$

$$\Delta P_{WTI,t} = a_{WTI} + \alpha_{WTI}\ \mu_{t-1} + \sum_{i=1}^{9} b_{WTI,i}\ \Delta P_{WTI,t-i} + \sum_{i=1}^{9} c_{WTI,i}\ \Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{WTI,j}\ X_{j,t} + \varepsilon_{WTI,t}\,, \tag{3}$$

Equation (1) is the cointegrating equation which captures the first component, the hypothesized fundamental tie between Henry Hub natural gas and WTI crude oil prices. $P_{HH,t}$ is the log natural gas price in week t, $P_{WTI,t}$ is the log crude oil price in week t, γ is a constant to be estimated, β is a parameter to be estimated, and μ_t is the error term in week t. $\Delta P_{HH,t}$ is the change in the log natural gas price from week t-1 to week t, μ_{t-1} is the lagged set of equilibrium errors from equation (1), $\Delta P_{WTI,t}$ is the change in the log crude oil price, X_j is the matrix of six exogenous variables representing additional drivers of the Henry Hub natural gas price, the variously subscripted parameters a, α, b, c, and d are to be estimated. Finally, $\varepsilon_{HH,t}$ and $\varepsilon_{WTI,t}$ are the error terms. After estimating the VECM, we then estimate the matching conditional Error Correction Model (conditional ECM),

$$\Delta P_{HH,t} = a_{ECM} + \alpha_{ECM}\ \mu_{t-1} + b_{ECM}\ \Delta P_{WTI,t} + \sum_{i=1}^{9} c_{ECM,i}\ \Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{ECM,j}\ X_{j,t} + \varepsilon_{ECM,t}. \quad (4)$$

Equation (4) uses as an input, μ_{t-1}, the previously estimated error term from equation (1) together with the contemporaneous change in the price of oil, the lagged changes in the price of natural gas, and the matrix of six exogenous variables. These capture the second and third components of the relationship between the natural gas and oil prices, and the error term captures the fourth component, the residual volatility.

The results of the estimation are shown in Tables 1 and 2. Following Brown and Yücel (2008), the exogenous variables used are the number of heating degree days (HDD), the number of cooling degree days (CDD), the deviation from the normal number of heating degree days (HDDDEV), the deviation from the normal number of cooling degree days (CDDDEV), the deviation of the amount of natural gas in storage from its average (STORDIFF), and the amount of natural gas production shut-in, e.g. due to storms (SHUTIN).[8] Because certain of the exogenous variables are only available starting June 13, 1997, when we report results for the cointegration analysis, the results are based on estimation over the June 13, 1997 to December 31, 2010 period of time, and do not include the period 1991 up to June 13, 1997.

---

[8] Normal Heating Degree Days or Cooling Degree Days reflect the average value for each week from 1971-2000.

We now turn to discussing the results in more detail, focusing one at a time on the separate components, beginning with the first.

The Fundamental Tie Between the Natural Gas Price and the Oil Price

The estimated cointegrating equation is:

$$P_{HH} = -0.0333 + (0.468 \times P_{WTI}). \tag{5}$$

This relationship is graphed in Figures 2 and 3. The cointegrating relationship is linear in the logged prices. Converted back into dollars, the log-linear relationship is a slightly concave curve. As can be seen in the Figures, when the oil price is $10/bbl, the cointegrating relationship predicts a natural gas price of $2.84/mmBtu. At $60/bbl, the predicted natural gas price is $6.57/mmBtu. Were the oil price to reach $150/bbl, the cointegrating relationship predicts a corresponding natural gas price of $10.09/mmBtu.

One can see in Figures 2 and 3 that the cointegrating equation attempts to fit the data better than the various rules-of-thumb by crafting a compromise out of slightly overestimating natural gas prices when the oil price is low and slightly underestimating them when the oil price is high. The fact that the equation is concave when converted out of the log-linear form in which it is estimated also makes the fit better. Nevertheless, it is impossible to escape the problem of the great volatility in natural gas prices. Figure 6 provides a time-series representation of the performance of the cointegrating relationship—equation (5)—at predicting the natural gas price from January 1991 through December 2010, i.e. both before and during the sample period used in the estimation. The centered mean absolute error for the cointegrating relationship is 0.394, which is approximately the same as for the burner tip distillate rule-of-thumb discussed above.[9] This repeats the earlier observation that the natural gas price series is just too volatile to be accounted for by any simple tie to the oil price including this cointegrating equation. Only by somehow accounting for this additional volatility could we reduce this error. The other components of the VECM and the conditional ECM attempt to provide this accounting, and we now turn to examine how successfully they do so.

The Error Correction Mechanism and the Rate of Recovery

Many factors may pull the price of natural gas away from the fundamental relationship. The model then allows for an error correction mechanism by which the natural gas price is pulled back to the fundamental relationship. This reversion to the fundamental relationship is a predictable part of the price movement captured in the estimated error correction mechanism. For example, when the crude oil price rises 20%, from $50/bbl to $60/bbl, and all other variables are held constant, then, according to the cointegrating relationship, the price of natural gas should rise approximately 9%, from $6.04/mmBtu to $6.56/mmBtu. This occurs gradually, however, with the half-life of the rise being nearly 22 weeks. Alternatively, if the price of natural gas price spikes up by 166%, from $6.04/mmBtu to $10/mmBtu, while the crude oil price is steady at $50/bbl, then the natural gas price is expected to eventually fall back to $6.04/mmBtu. The half-life for the return of the natural gas price to its cointegrating relationship is nearly 8 weeks.

[9] Note that the VECM was estimated over the period June 13, 1997-December 31, 2010, so the 0.394 mean absolute error incorporates errors both in- and out-of-sample. Focusing just on the 1997-2010 data used for the estimation, the centered mean absolute error for the cointegrating relationship is 0.341. Over this shorter window, the centered mean absolute error for the rules-of-thumb ranged from 0.398 for the 10-to-1 and the energy content equivalence rules to 0.411 for the residual burner tip parity rule, to 0.428 for the distillate burner tip parity rule. So, not surprisingly, the cointegrating equation does fit the data in-sample better than any of the rules-of-thumb.

Exogenous Factors

The first two of our exogenous variables, HDD and CDD, capture the well known seasonality of natural gas prices in the U.S. To generate the estimated seasonal fluctuations, we simulate the path of natural gas prices through the average annual cycle of HDD and CDD in our dataset using our estimated conditional ECM and holding the crude oil price, the other four exogenous factors and the error terms all fixed. The resulting natural gas price settles into a cycle around an average point that occurs in the first week of July and again in the third week of December. The price peaks at about 113% in April. The trough is at 87% in September. The total amplitude of the seasonal variation in the natural gas price is 26 percentage points. At a base price of $7/mmBtu (the July and December prices), this is a range of $1.82/mmBtu. Figure 6 shows this seasonal variation overlayed on the observed prediction errors for the estimated cointegrating relationship. This allows one to see how much larger is the actual range of variation than can be accounted for by the predictable seasonal component. For example, the standard deviation of the logged error series for the cointegrating relationship from 1991–2010 is 0.394. The seasonality coefficient, however, only ranges as high or low as +/– 0.131, or less than half of a standard deviation. Using two standard deviations as a benchmark for capturing the vast majority of the range in gas volatility, the seasonal component could not account for any more than 16.6% of natural gas volatility.

The next two of our exogenous variables, HDDDEV and CDDDEV, capture the impact of unseasonably cold or warm weather on demand and therefore price. These variables, too, only account for a modest amount of the volatility in the natural gas price. To illustrate this, Figure 7 compares the actual changes in the price of natural gas prices around a typical cold spell—the two weeks of March 18-25, 2005—against the portion of the price change attributable to the cold spell. A typical spell last two weeks, with the first week exhibiting an HDD level 20 degree days above normal and the second week exhibiting an HDD level 12 degree days above normal. The portion attributable to the cold spell is calculated using the estimated conditional ECM in equation (4) using the actual deviation in HDD, and holding the crude oil price fixed, setting the initial natural gas price so that the error term in equation (1) is zero, setting the other exogenous factors in equation (4) to zero, setting the error terms in equation (4) to zero, and, simulating how the natural gas price evolves in response to the shock to the exogenous variable HDDDEV. In our example, the price on week 0, March 11, 2005, was $7.51/mmBtu. On March 18, week 1 of the cold spell, the price had fallen by $0.05/mmBtu. The cold spell is predicted to have increased the price by $0.13/mmBtu, so that absent the cold spell the price would have fallen by $0.18/mmBtu. By March 18, week 2 of the cold spell, the price rose sharply yielding a cumulative increase of $0.50/mmBtu. The estimated model attributes a cumulative increase of $0.19 to the cold spell, accounting for about half of the actual cumulative increase. From there on out, the cumulative increase attributable to the cold spell gradually dissipates. The cumulative change in the actual price swings far below and far above zero.

There are two other exogenous variables – the level of natural gas storage and the level of shut-in production in the Gulf of Mexico due to hurricanes. These, too, account for a modest amount of the volatility of natural gas prices, although for economy of space we provide detail on the impact of these two individual variables in the Appendix.

The Residual or Unexplained Volatility

The implementation of the VECM and conditional ECM modeling techniques improves the fit of the predicted natural gas price over the rules-of-thumb. Nevertheless, a large amount of the volatility in natural gas prices could not be explained by the combination of the cointegrating relationship with the crude oil price, the error correction mechanism, and the identified exogenous variables. The portion of the volatility in the natural gas price explained by the conditional ECM is approximately 15%. That means the fraction of variance in natural gas prices unexplained by our model is nearly 85%. Therefore, although the two series are cointegrated, this statistical fact should not be taken to mean that the two series are tightly coupled. Over short horizons there is significant unexplained volatility in the natural gas price. The two prices regularly decouple, sometimes significantly, although this decoupling is not long lasting.

4. A CHANGING RELATIONSHIP OVER TIME?

One possible explanation for the weak explanatory power of the model is that we are trying to identify a single relationship across a long window of time, when in fact the relationship has evolved over this period. As we noted earlier, the natural gas and oil prices are not likely to be equated simply on the basis of energy equivalence because of the different costs of production, transportation, processing and storage, and because of the different end use markets they serve. These different underlying technical and economic factors make the equilibrium relationship diverge from a strict energy equivalence. But these factors are themselves shifting over time, sometimes gradually and sometimes swiftly. Villar and Joutz (2006) examined the 1989-2005 period and found that the cointegrating relationship between logged oil and gas prices shifted up by nearly half of a percent per month, with the price of natural gas relative to crude oil having increased. Hartley, Medlock and Rosthal (2008) examined a substantially overlapping period, 1990-2006, which exhibited a similar increase in the price of natural gas relative to the price of crude oil. Hartley, Medlock and Rosthal go a step further to specifically attribute this to the increased demand for natural gas arising from the installation of advanced combined cycle gas turbine (CCGT) power plants with significantly improved heat rates. Figure 8 shows the dramatic shift up in the cointegrating relationship documented by Villar and Joutz, contrasting the estimated relationship at the start of their data set, in 1989, with the estimated relationship at the end, in 2005.

Has the cointegrating relationship shifted once again, but this time in the opposite direction? A major technological innovation in recent years has been improvements in the horizontal drilling and hydraulic fracturing making possible the low cost exploitation of natural gas in shales. Production from shales has dramatically increased in the U.S. in recent years, and is almost certainly the cause of the most recent drop in the price of natural gas relative to oil. Simultaneously, the price of oil has reached a higher level than before, and oil use is more and more dominated by the transport sector. Each of these developments shapes the competition between the two energy carriers and therefore the equilibrium relationship between them. Is this shift statistically identifiable in our data, taking into account the error correction mechanism and exogenous factors?

To address this question, we examined our data as follows. First, as reported in the previous section, we fit a single cointegrating relationship over our full dataset, June 13, 1997 to December 31, 2010.[10]

Second, we considered the possibility of cointegration, but allowing for breakpoints in the structure of the relationship. We identify two breakpoints, one at February 6, 2009 and one at March 10, 2006.[11]

Looking at the data up to February 6, 2009, the evidence for a cointegrating relationship is strong. This is true whether we fit a single relationship across the full window of time, June 13, 1997-February 6, 2009, or we fit two separate relationships for the two sub-segments, June 13, 1997-March 10, 2006 and March 17, 2006-February 6, 2009. However, the evidence clearly argues that the relationship shifted across the two sub-segments.[12] The two cointegrating relationships we estimate for our two windows of time are:

$$log\ P_{HH} = -1.2007 + (0.7261 \text{ x } log\ P_{WTI}) \quad (6)$$

---

[10] We ran the Johansen test for cointegration over our full window of time, from June 13, 1997 through December 31, 2010, with lag length of 10 weeks. The test selects a rank of 1 over a rank of 2 at the 1% significance level, but also selects a rank of 0 (not cointegrated) over a rank of 1 (cointegrated) at the 5% level. The Schwartz-Bayesian Information Criterion and the Hannan-Quinn Information criteria disagree on the rank, with the former selecting a rank of 0 and the latter selecting a rank of 1. If we select a rank of 1 (cointegrated) and fit our VECM over the full window of time and then evaluate the errors for a unit root using the Augmented Dickey Fuller tests or the Phillips-Perron tests, we reject the null hypothesis of a unit root in the errors, which is evidence that the identified relationship is a cointegrating relationship.

[11] We employ the Gregory and Hansen (1996) tests for cointegration with regime shift. The null hypothesis is no cointegration across the full time period, and the alternative hypothesis is cointegration, where the cointegrating vector is allowed to change at a single unknown time during the sample period. The alternative includes the possibility of no change or break. Their tests do not require ex ante information on the timing of a break, nor a presumption about whether or not there is a break. All three tests accept the alternative of cointegration, including the possibility of a breakpoint. Although there is minor disagreement among the tests about the exact dating of a likely breakpoint, we chose February 6, 2009 based on the ADF test. We then repeated the Gregory and Hansen (1996) test on the shorter window of time from June 13, 1997 through February 6, 2009 and identified the earlier breakpoint at March 10, 2006.

[12] Having chosen a specific break point, it is appropriate to apply the Chow test to determine stability of the estimated intercept and slope coefficient in the cointegrating relationship. This establishes that the values are not constant across the two periods.

for the June 13, 1997-March 10, 2006 period, and

$$log\ P_{HH} = 0.1969 + (0.4621 \text{ x } log\ P_{WTI}) \tag{7}$$

for the March 17, 2006-February 6, 2009 period. Figure 9 graphs these two relationships. The cointegrating equation has shifted downward in the latter period, predicting a lower price of natural gas given the price of crude oil. This shift is in exactly the opposite direction from the shift documented by previous authors for the earlier era, 1989-2005.

Finally, looking at the short window of time from February 13, 2009 through December 31, 2010, we cannot say much since neither series displays sufficient evidence of non-stationarity for the tests of cointegration to be meaningful. The Appendix contains a description of the full set of tests and investigations performed across all segments and combinations.

Of course, the conditional error correction models based on each of the segmented cointegrating relationships also account for a greater portion of the volatility in natural gas than the model covering the period as a whole. The June 13, 1997-March 10, 2006 model accounts for nearly 21% of natural gas volatility through the crude oil price and the included conditioning variables. The March 17, 2006-February 6, 2009 model accounts for 26% of the price volatility in natural gas. Nevertheless, there remains a large amount of unexplained volatility in the natural gas price even in each of these separately estimated time windows.

These results support the hypothesis that whatever relationship might characterize the prices of natural gas and oil, that relationship is not stable over long periods of time. Earlier researchers documented a statistically reliable relationship through a window of years when the price of gas shifted upward relative to the price of oil, and we have documented a statistically reliable relationship during subsequent years when the price of gas was lower relative to the price of oil. Today's tie between the price of natural gas and the price of oil may not be very predictive of tomorrow's tie.

CONCLUSION

A number of recent academic studies have established that natural gas and crude oil prices are cointegrated. However, recent years have witnessed a price of natural gas that seems decoupled from the price of oil, reaching new lows relative to the price of oil. In this paper we have confronted the apparent contradiction between these two facts by examining more closely what is and is not established by the cointegration tests. While we are able to reconfirm the presence of a statistically significant relationship between the two price series, our results emphasize two other points that are important to any discussion about a relationship or a decoupling.

First, there is an enormous amount of unexplained volatility in natural gas prices. The raw price series for natural gas is approximately twice as volatile as the raw oil price series. Applying a VECM and estimating a conditional ECM to account for the predictable error correction and for exogenous variables which temporarily disturb the relationship still leaves an enormous amount of volatility in the natural gas price unaccounted for. Our model of the 1997-2010 period only accounts for about 15% of the volatility in natural gas prices, leaving 85% unaccounted for. Splitting the sample up into shorter periods produces only a modest improvement in the fit, in-sample. There is no escaping the significant size of the short-run swings in the natural gas price that cannot be accounted for. At short horizons, the cointegrating relationship is statistically identified, but not very reliable for predicting the natural gas price with any precision.

Second, the cointegrating relationship itself has changed over time, shifting upward in one era and downward again in a later era. These shifts are likely due to shifts in the underlying technological and economic forces determining an equilibrium relationship between the two prices. Therefore, the historical cointegrating relationship may not be a very reliable predictor of the future natural gas price, at least not at longer horizons over which shifts in the underlying forces are unpredictable.

This analysis can inform the repeated discussions about how the natural gas price has “decoupled” from the oil price. First, our documentation of the unaccounted for volatility points out that there are likely to be many occasions when the prices temporarily break away from the usual relationship to which they will later return. These decouplings can be severe, but they are also not very long lasting – less than one season typically – and the old relationship is reestablished. Second, our documentation that the cointegrating relationship has shifted over time, first in one direction and then in another, points out that prices can decouple from one relationship only to recouple in a new relationship. Third, there is not yet any evidence that the relationship between the two price series has been severed completely. Indeed, it is hard to imagine that natural gas and oil prices could decouple completely and permanently. For example, while conversion of gas to liquids may seem expensive now, the technological possibility of conversion does place a cap on the degree to which oil prices can rise relative to natural gas prices. Other technological and economic constraints act similarly to prevent a complete decoupling. However, the freedom of motion is large.

REFERENCES

Adelman, M.A., and Watkins, G.C., 1997, The value of United States oil and gas reserves, Advances in the Economics of Energy and Resources, vol. 10, JAI Press Inc., pp. 131-184.

Brown, Stephen P. A., and Mine K. Yücel, 2008, What drives natural gas prices, Energy Journal 29(2), 45-60.

Gregory, Allan W., and Bruce E. Hansen, 1996, Residual-based tests for cointegration in models with regime shifts, Journal of Econometrics 70, 99-126.

Hartley, Peter R., Kenneth B. Medlock III, and Jennifer E. Rosthal, 2008, The relationship of natural gas to oil prices, Energy Journal 29(3), 47-65.

Hendry, David F., and Katarina Juselius, 2000, Explaining cointegration analysis: Part 1, Energy Journal 21(1), 1-42, and Part 2, Energy Journal 22(1), 75-120.

IEA (International Energy Agency), 2009, World Energy Outlook. OECD. Paris, France.

Serletis, Apostolos and John Herbert, 1999, The message in North American energy prices, Energy Economics 21, 471-483.

Smith, J.L., 2004, Petroleum property valuation, Encyclopedia of Energy, vol. 4, Elsevier, pp. 815-816.

Villar, Jose, A., and Frederick L. Joutz, 2006, The relationship between crude oil and natural gas prices, Energy Information Administration, Office of Oil and Gas.

Table 1
Parameter Estimates for the Vector Error Correction Model
Full Period, June 13, 1997-December 31, 2010

Variable	Coefficient	Estimate	p-value	
	γ	-0.0333		
$P_{WTI\ t}$	β_{WTI}	0.4680	0.001	**

Variable	Coefficient	Estimate	p-value		Variable	Coefficient	Estimate	p-value	
	a_{HH}	-0.0018	0.902			a_{WTI}	-0.0076	0.369	
μ_{t-1}	α_{HH}	-0.0468	0.001	**	μ_{t-1}	α_{WTI}	0.0109	0.203	
$\Delta P_{WTI\ t-1}$	$b_{HH\ 1}$	0.0843	0.199		$\Delta P_{WTI\ t-1}$	$b_{WTI\ 1}$	-0.1124	0.004	**
$\Delta P_{WTI\ t-2}$	$b_{HH\ 2}$	-0.0146	0.824		$\Delta P_{WTI\ t-2}$	$b_{WTI\ 2}$	-0.0959	0.014	*
$\Delta P_{WTI\ t-3}$	$b_{HH\ 3}$	0.0647	0.327		$\Delta P_{WTI\ t-3}$	$b_{WTI\ 3}$	0.0837	0.032	*
$\Delta P_{WTI\ t-4}$	$b_{HH\ 4}$	-0.0726	0.270		$\Delta P_{WTI\ t-4}$	$b_{WTI\ 4}$	-0.0322	0.409	
$\Delta P_{WTI\ t-5}$	$b_{HH\ 5}$	-0.0320	0.627		$\Delta P_{WTI\ t-5}$	$b_{WTI\ 5}$	0.0207	0.596	
$\Delta P_{WTI\ t-6}$	$b_{HH\ 6}$	0.0660	0.317		$\Delta P_{WTI\ t-6}$	$b_{WTI\ 6}$	-0.0497	0.202	
$\Delta P_{WTI\ t-7}$	$b_{HH\ 7}$	0.1793	0.007	**	$\Delta P_{WTI\ t-7}$	$b_{WTI\ 7}$	-0.0711	0.068	+
$\Delta P_{WTI\ t-8}$	$b_{HH\ 8}$	0.0744	0.260		$\Delta P_{WTI\ t-8}$	$b_{WTI\ 8}$	0.1331	0.001	**
$\Delta P_{WTI\ t-9}$	$b_{HH\ 9}$	0.0833	0.206		$\Delta P_{WTI\ t-9}$	$b_{WTI\ 9}$	0.0969	0.013	*
$\Delta P_{HH\ t-1}$	$c_{HH\ 1}$	-0.1057	0.008	**	$\Delta P_{HH\ t-1}$	$c_{WTI\ 1}$	0.0547	0.020	*
$\Delta P_{HH\ t-2}$	$c_{HH\ 2}$	-0.0548	0.163		$\Delta P_{HH\ t-2}$	$c_{WTI\ 2}$	0.0076	0.743	
$\Delta P_{HH\ t-3}$	$c_{HH\ 3}$	-0.1498	0.000	**	$\Delta P_{HH\ t-3}$	$c_{WTI\ 3}$	-0.0143	0.530	
$\Delta P_{HH\ t-4}$	$c_{HH\ 4}$	-0.0414	0.283		$\Delta P_{HH\ t-4}$	$c_{WTI\ 4}$	0.0142	0.533	
$\Delta P_{HH\ t-5}$	$c_{HH\ 5}$	-0.1005	0.008	**	$\Delta P_{HH\ t-5}$	$c_{WTI\ 5}$	0.0162	0.473	
$\Delta P_{HH\ t-6}$	$c_{HH\ 6}$	0.0498	0.193		$\Delta P_{HH\ t-6}$	$c_{WTI\ 6}$	0.0208	0.359	
$\Delta P_{HH\ t-7}$	$c_{HH\ 7}$	-0.0631	0.096	+	$\Delta P_{HH\ t-7}$	$c_{WTI\ 7}$	-0.0045	0.842	
$\Delta P_{HH\ t-8}$	$c_{HH\ 8}$	0.0184	0.628		$\Delta P_{HH\ t-8}$	$c_{WTI\ 8}$	-0.0153	0.494	
$\Delta P_{HH\ t-9}$	$c_{HH\ 9}$	-0.0832	0.028	*	$\Delta P_{HH\ t-9}$	$c_{WTI\ 9}$	-0.0309	0.167	
HDD_t	$d_{HH\ 1}$	8.58E-05	0.311		HDD_t	$d_{WTI\ 1}$	5.46E-05	0.276	
$HDDev_t$	$d_{HH\ 2}$	1.03E-03	0.000	**	$HDDev_t$	$d_{WTI\ 2}$	-7.57E-05	0.528	
CDD_t	$d_{HH\ 3}$	-4.68E-04	0.072	+	CDD_t	$d_{WTI\ 3}$	1.43E-04	0.352	
$CDDev_t$	$d_{HH\ 4}$	3.35E-03	0.000	**	$CDDev_t$	$d_{WTI\ 4}$	3.92E-04	0.339	
$StorDiff_t$	$d_{HH\ 5}$	-1.94E-05	0.252		$StorDiff_t$	$d_{WTI\ 5}$	1.88E-05	0.060	+
$Shutin_t$	$d_{HH\ 6}$	4.55E-06	0.290		$Shutin_t$	$d_{WTI\ 6}$	-6.25E-06	0.014	*

Variables	Chi^2 Stat	p-value		Variables	Chi^2 Stat	p-value	
Lagged HH	45.67	0.000	**	Lagged HH	10.38	0.321	
Lagged WTI	13.57	0.138		Lagged WTI	49.84	0.000	**
Lagged HH + WTI	59.68	0.000	**	Lagged HH + WTI	56.29	0.000	**
Exogenous Variables	50.77	0.000	**	Exogenous Variables	12.43	0.053	+
Exog + Lagged HH	86.03	0.000	**	Exog + Lagged HH	24.23	0.061	+
Exog + Lagged WTI	61.09	0.000	**	Exog + Lagged WTI	62.07	0.000	**
Exogenous + Lagged	98.91	0.000	**	Exogenous + Lagged	69.94	0.000	**

Equation	parameters	RMSE	R-sq	Chi^2 Stat	p-value
ΔP_{HH}	26	0.0958	0.1458	114.7092	0.0000
ΔP_{WTI}	26	0.0567	0.0952	70.6978	0.0000

+ = 0.1, * = 0.05, ** = 0.01 significance levels. Number of Observations: 698

Table 2
Parameter Estimates for the conditional Error Correction Model
Full Period, June 13, 1997-December 31, 2010

Variable	Coefficient	Estimate	p-value	
	a_{ECM}	0.0044	0.750	
μ_{t-1}	α_{ECM}	-0.0497	0.000	**
$\Delta P_{WTI\ t}$	b_{ECM}	0.2458	0.000	**
$\Delta P_{HH\ t-1}$	$c_{ECM\ 1}$	-0.1040	0.007	**
$\Delta P_{HH\ t-2}$	$c_{ECM\ 2}$	-0.0463	0.216	
$\Delta P_{HH\ t-3}$	$c_{ECM\ 3}$	-0.1424	0.000	**
$\Delta P_{HH\ t-4}$	$c_{ECM\ 4}$	-0.0472	0.200	
$\Delta P_{HH\ t-5}$	$c_{ECM\ 5}$	-0.1090	0.003	**
$\Delta P_{HH\ t-6}$	$c_{ECM\ 6}$	0.0563	0.125	
$\Delta P_{HH\ t-7}$	$c_{ECM\ 7}$	-0.0393	0.279	+
$\Delta P_{HH\ t-8}$	$c_{ECM\ 8}$	0.0376	0.299	
$\Delta P_{HH\ t-9}$	$c_{ECM\ 9}$	-0.0686	0.059	*
HDD_t	$d_{ECM\ 1}$	3.91E-05	0.630	
$HDDev_t$	$d_{ECM\ 2}$	1.06E-03	0.000	**
CDD_t	$d_{ECM\ 3}$	-5.29E-04	0.037	+
$CDDev_t$	$d_{ECM\ 4}$	3.18E-03	0.000	**
$StorDiff_t$	$d_{ECM\ 5}$	-1.94E-05	0.239	
$Shutin_t$	$d_{ECM\ 6}$	4.40E-06	0.289	

Variables	Chi2 Stat	p-value	
Lagged HH	48.32	0.000	**
Exogenous Variables	50.35	0.000	**
Exog + WTI	65.37	0.000	**
Exog + Lagged HH	89.70	0.000	**
Exogenous + WTI + Lagge	104.74	0.000	**

Equation	parameters	RMSE	R-sq	Chi^2 Stat	p-value
ΔP_{HH}	18	0.0951	0.1479	121.1898	0.0000

+ = 0.1, * = 0.05, ** = 0.01 significance levels. Number of Observations: 698

Figure 1. The Natural Gas and Crude Oil Spot Prices, 1991-2010 (real 2010 dollars).

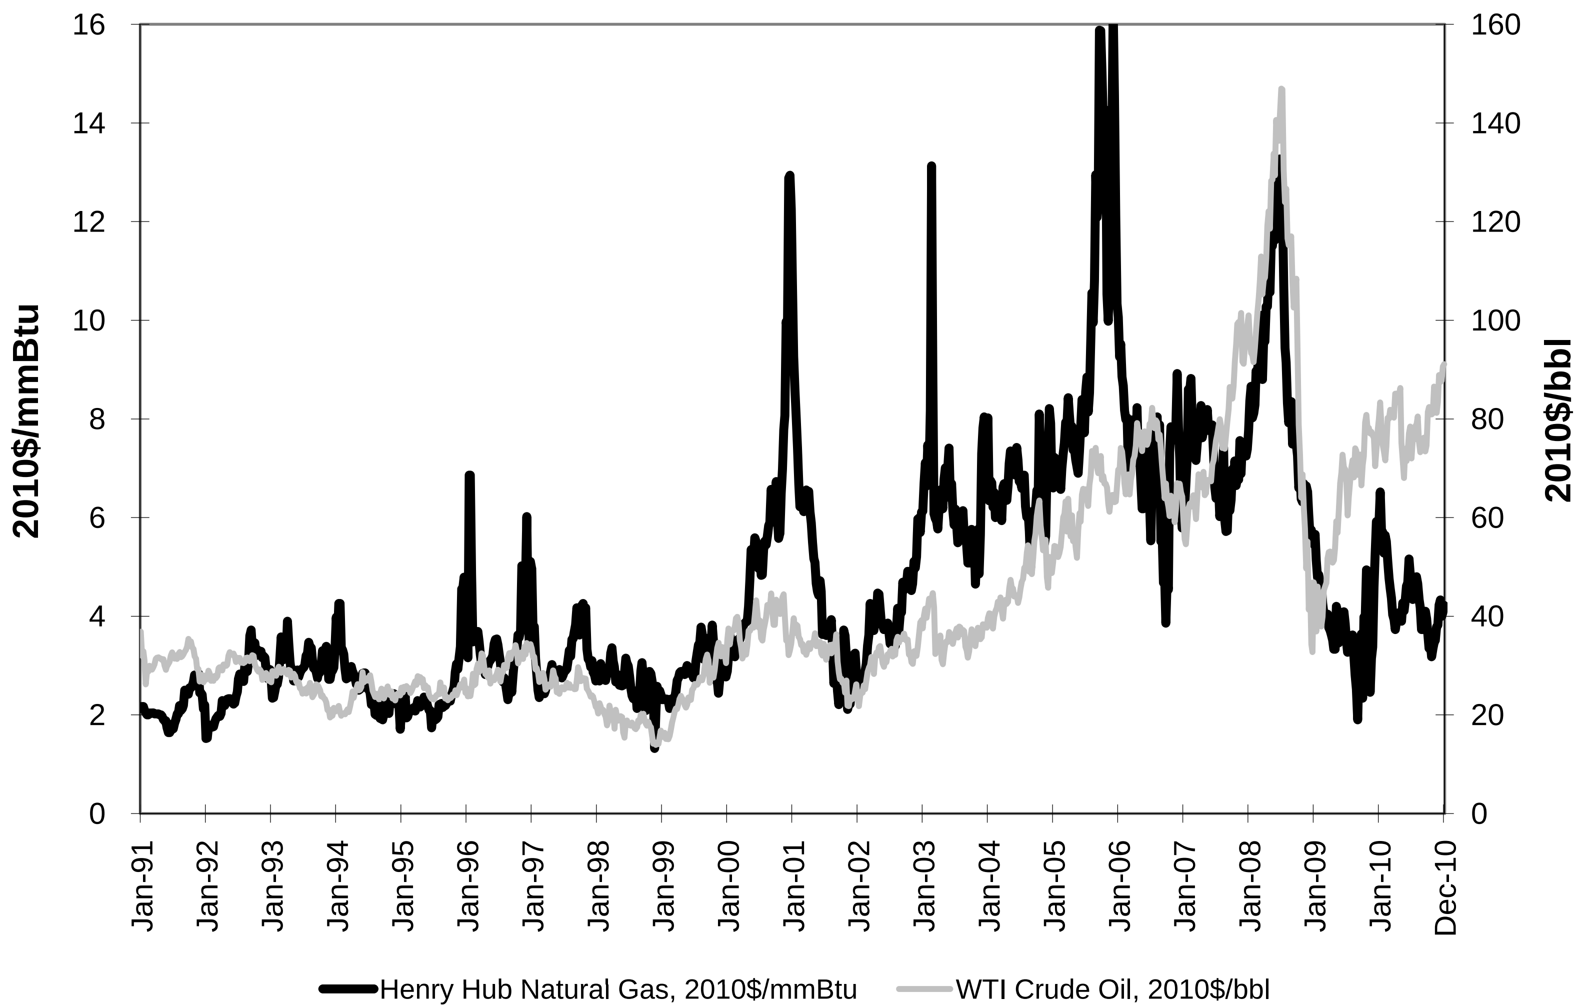

Both series are weekly day-ahead prices of commodities as sampled by Bloomberg. The natural gas prices are volume-weighted averages in $/mmBtu for delivery at Henry Hub in Louisiana. The crude oil prices are the arithmetic averages in $/bbl for West Texas Intermediate (WTI) crude oil traded at Cushing, Oklahoma. All prices were subsequently converted into real 2010 dollars.

Figure 2. Pricing Rules-of-Thumb Versus Observed Prices, 1991-2010

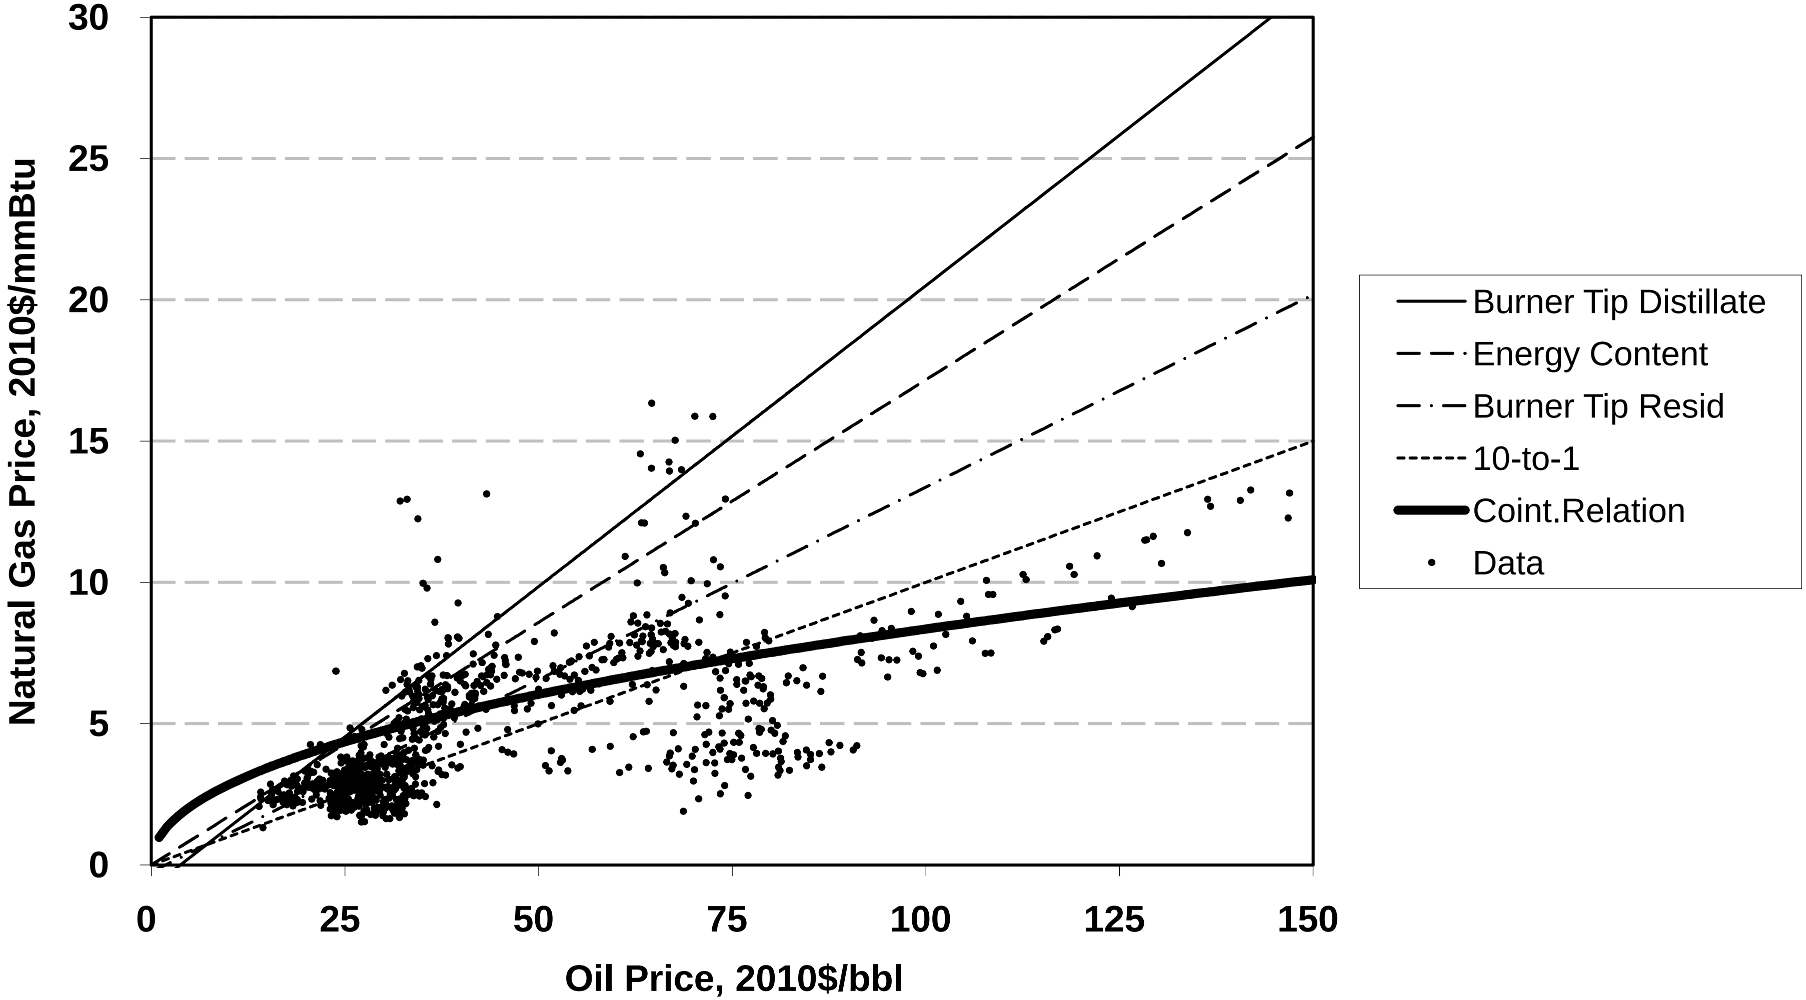

The figure charts natural gas prices as a function of oil prices. The four straight lines show the four pricing rules-of-thumb. The top line (using the ordering of the lines at the right of the figure) is the burner-tip parity rule based on natural gas competing with distillate fuel oil, the second line is the energy-content equivalence rule, the third line is the burner-tip parity rule based on natural gas competing with residual fuel oil, and the fourth line is the 10-to-1 rule. The dark black, slightly curved line is the estimated cointegrating equation from the VECM. The scatterplot of data points are the actual price combinations observed over the 1991-2010 period. All observed prices are quoted in real terms in 2010 dollars.

Figure 3. Pricing Rules-of-Thumb Versus Observed Prices (Low Oil Price Range)

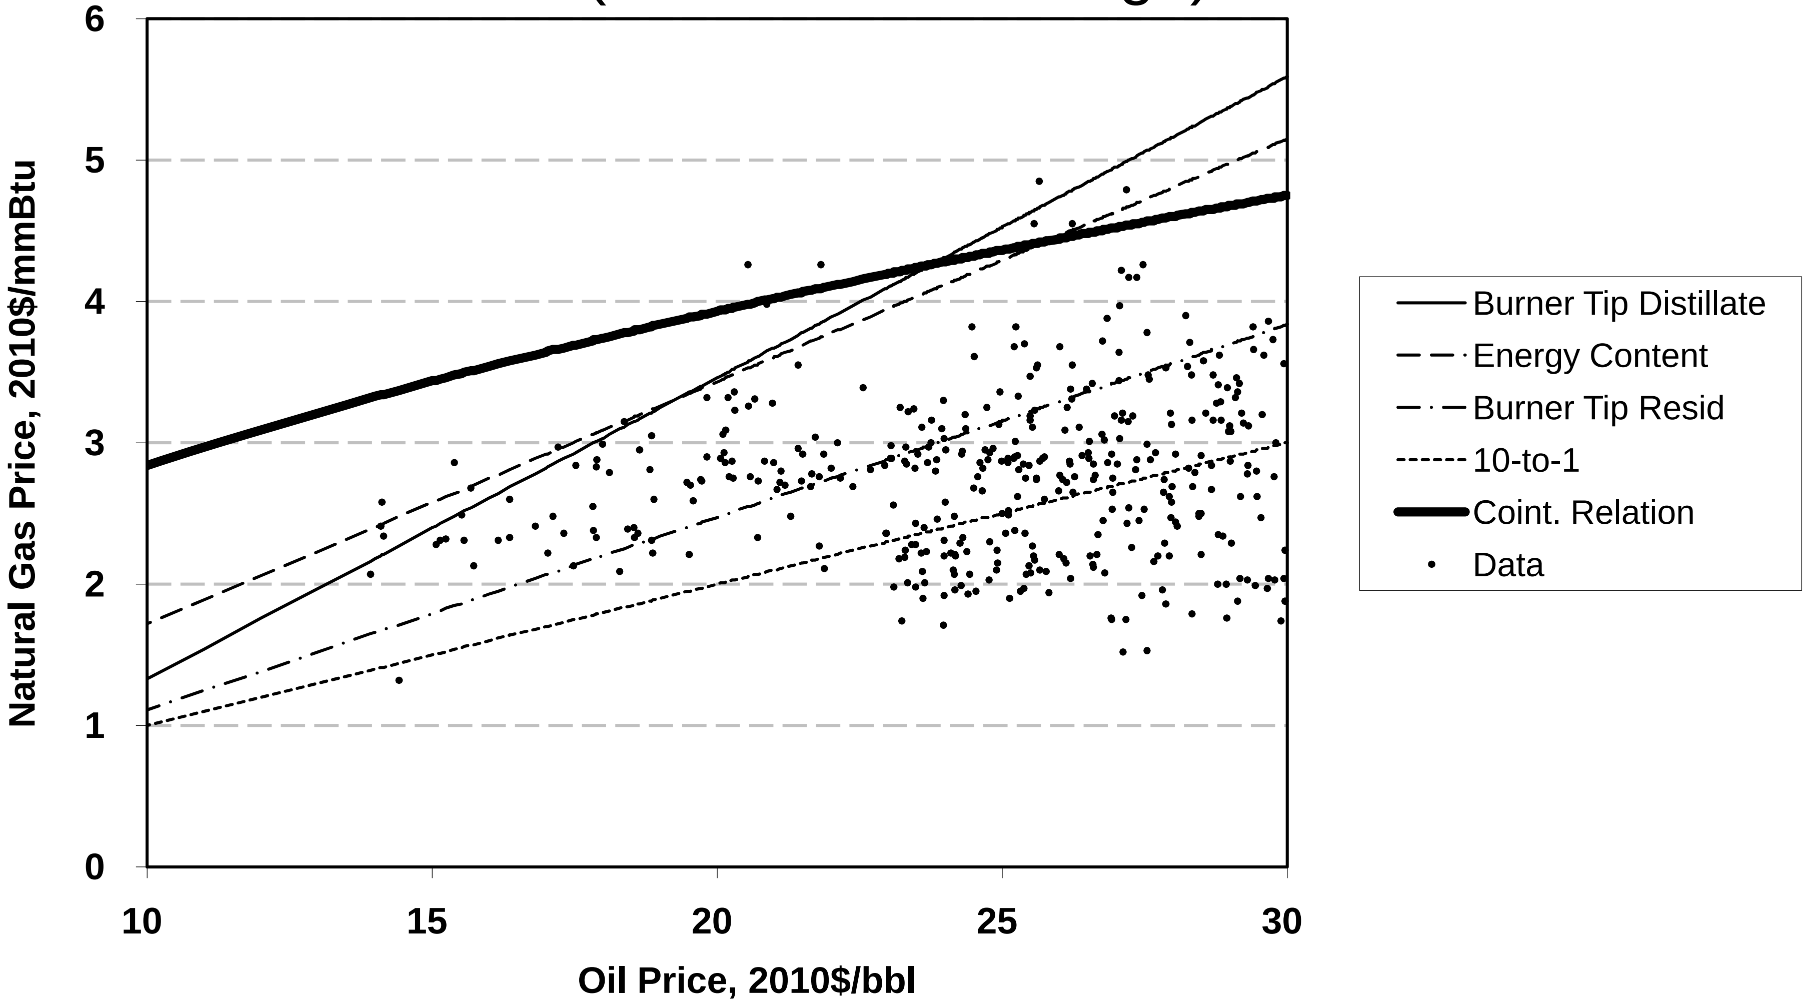

The figure shows the same data as Figure 2, except that it focuses in on the low range of oil prices so as to make visible the different observed prices and the comparison to the different rules-of-thumb.

Figure 4. Prediction Errors for 4 Rules of Thumb
(Actual Log Natural Gas Price minus Predicted Log Natural Gas Price)

Fig. 4A: Burner Tip Parity Rule: Distillate

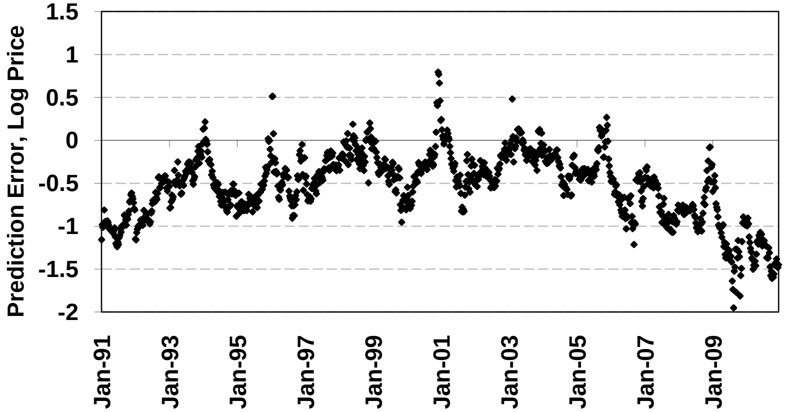

Fig. 4B: Energy Content Equivalence Rule

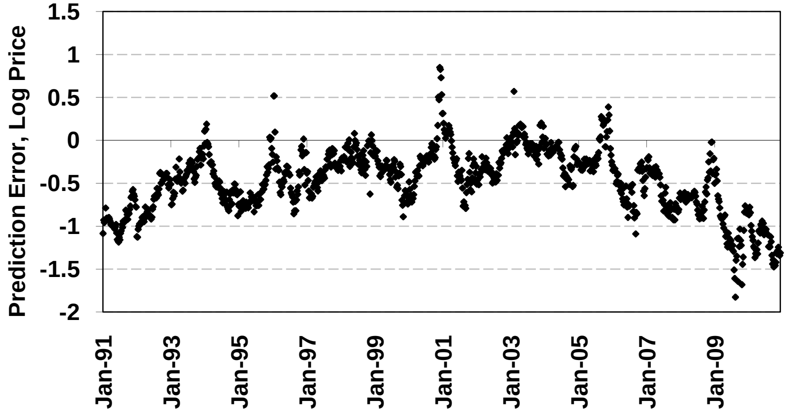

Fig. 4C: Burner Tip Parity Rule: Residual

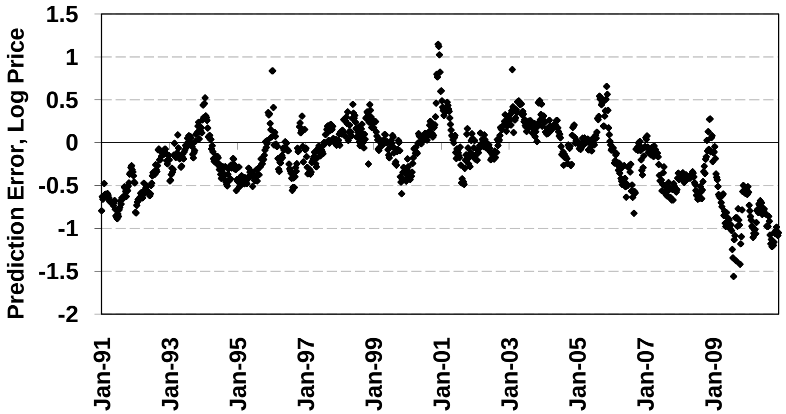

Fig. 4D: 10-to-1 Rule

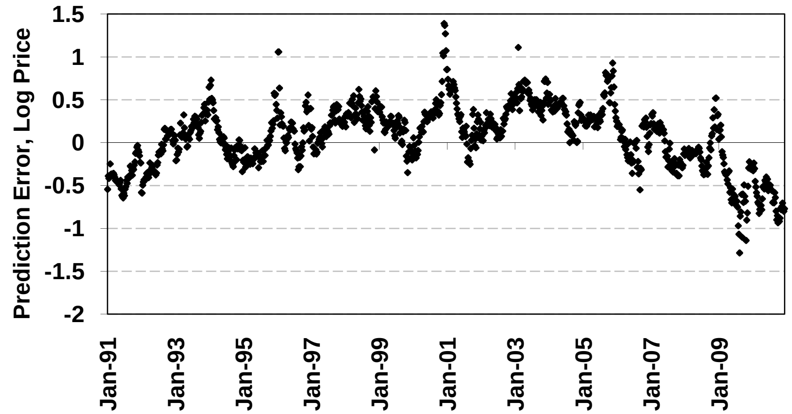

Figure 5. Prediction Errors for the Stand-alone Cointegrating Relationship

(Actual Log Natural Gas Price minus Long-Run Predicted Log Natural Gas Price)

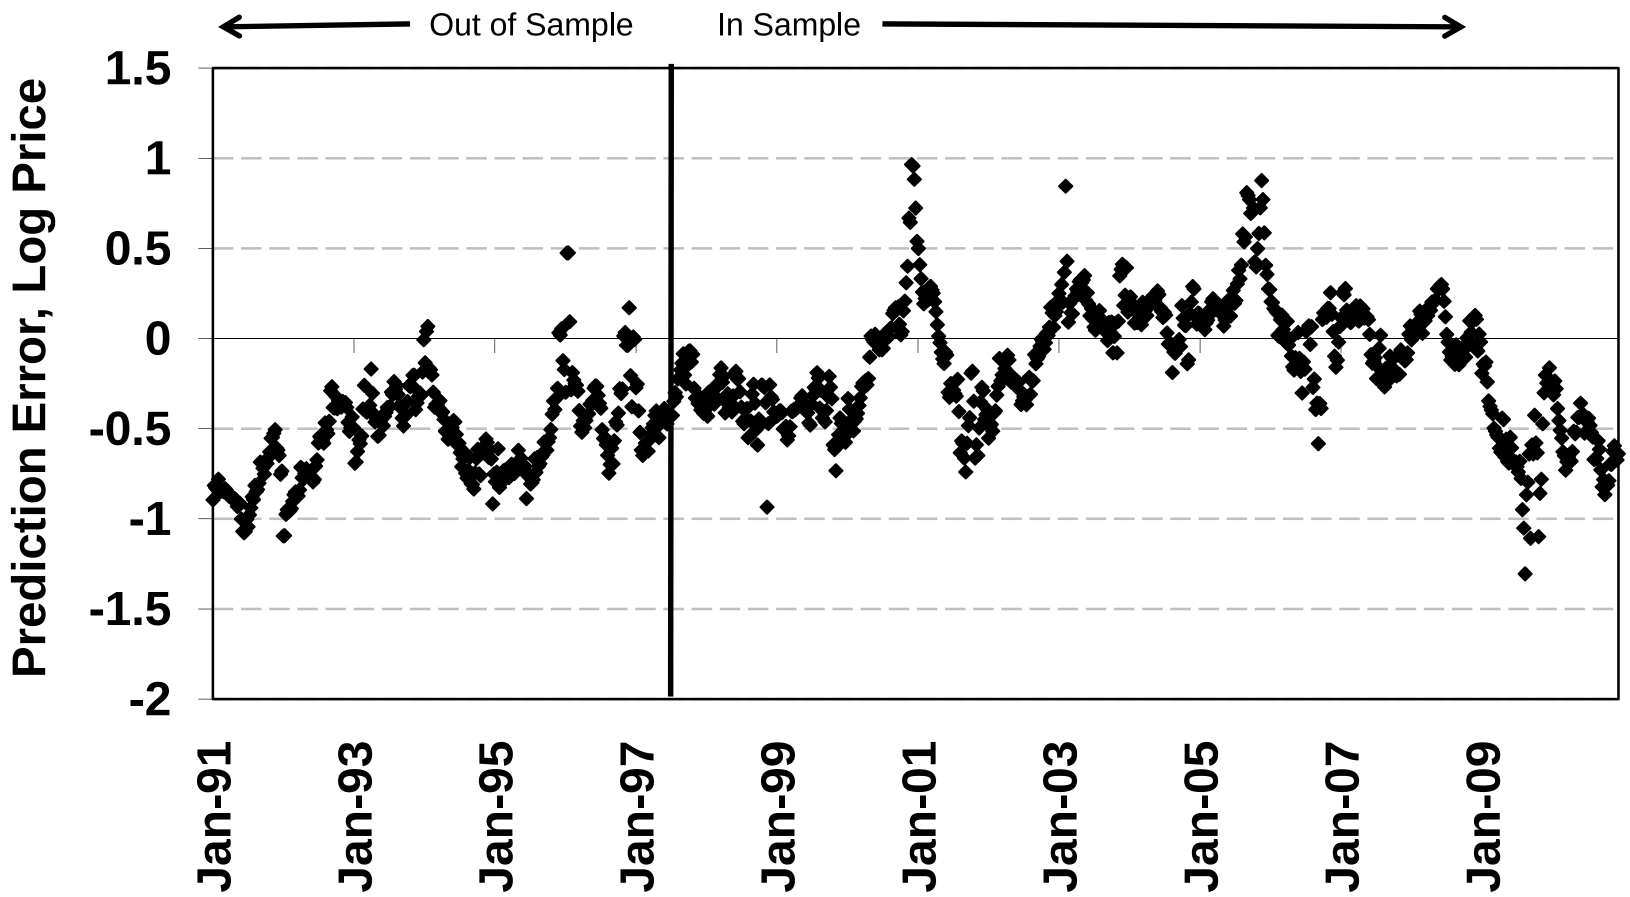

Prediction errors here are the difference between the actual natural gas price and the predicted natural gas price using equation (5).

Figure 6. Seasonality Relative to Prediction Errors

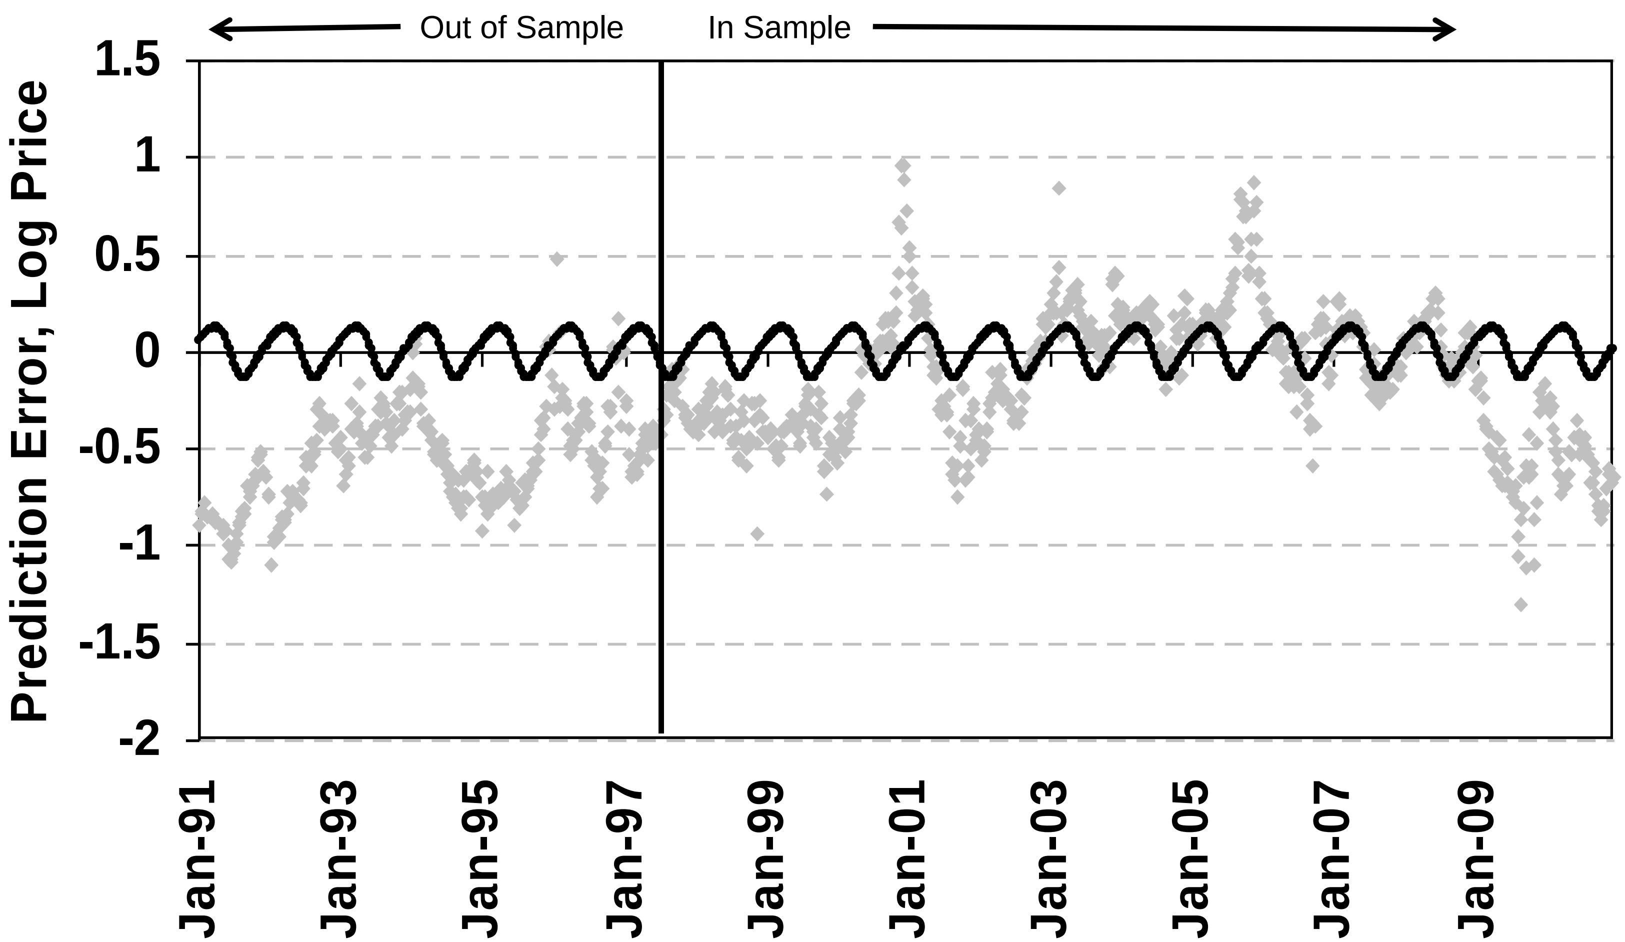

Figure 7. Predicted Impact of Unseasonal Cold Snap Versus Actual Change in Price.

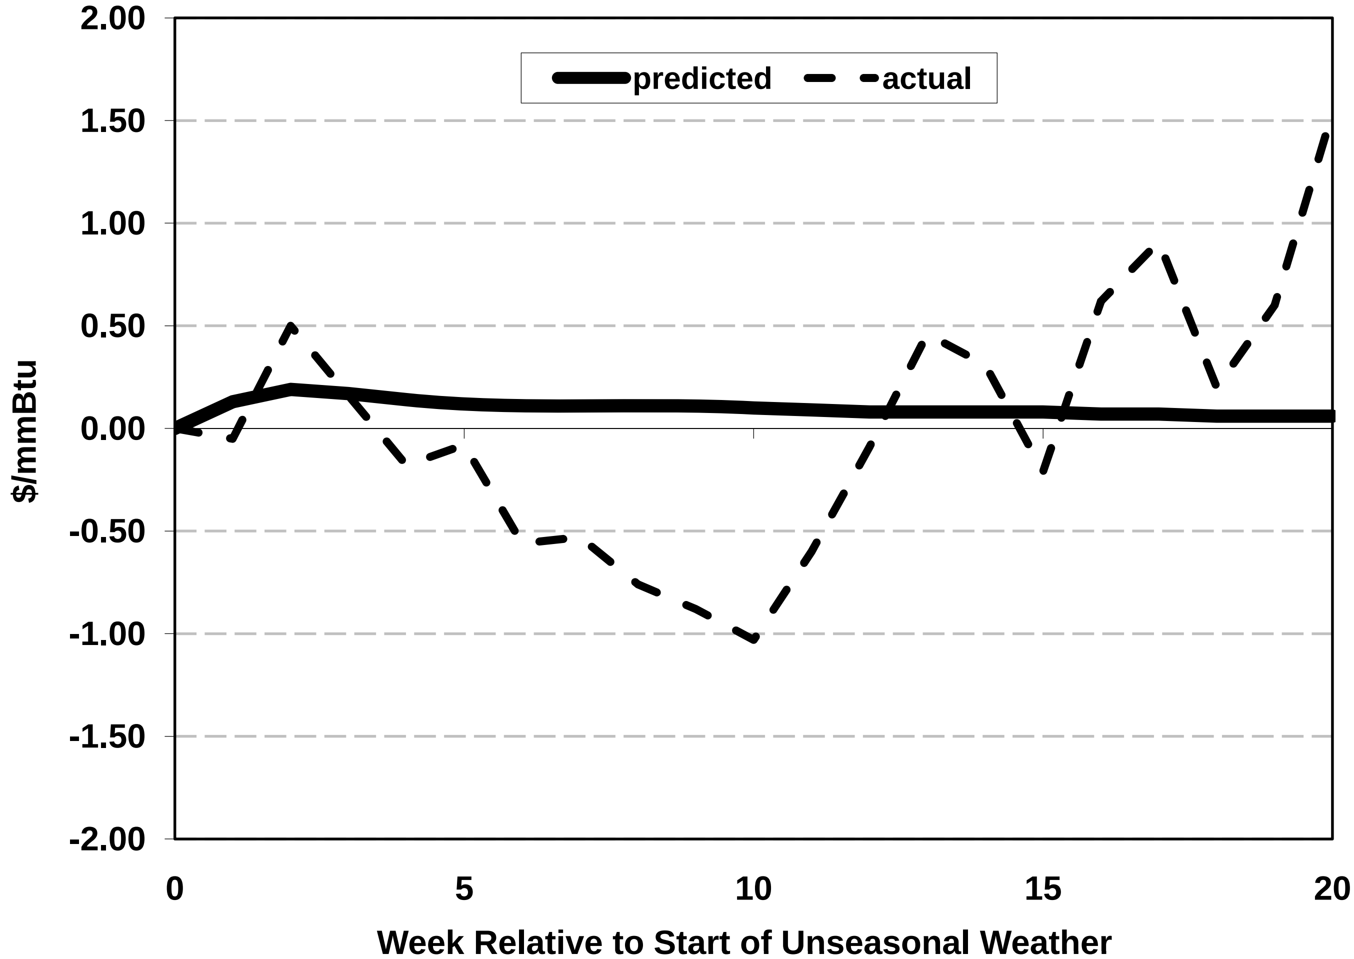

The figure shows the actual cumulative price change around a typical unseasonal cold snap. Week 0 is March 11, 2005, and weeks 1 and 2, March 18 and March 25, 2005 are the weeks of the cold snap. A typical spell last two weeks, with the first week exhibiting an HDD level 20 degree days above the average and the second week exhibiting an HDD level 12 degree days above the average. We chose this set of dates because it matched a typical cold snap. The figure shows also shows the predicted cumulative price impact of the cold spell. This is calculated using the estimated conditional ECM in equation (4) using the actual deviation in HDD, and holding the crude oil price fixed, setting the initial natural gas price so that the error term in equation (1) is zero, setting the other exogenous factors in equation (4) to zero, setting the error terms in equation (4) to zero, and, simulating how the natural gas price evolves in response to the shock to the exogenous variable of HDD deviations.

Figure 8. Villar-Joutz Shifting Cointegrating Relationship, 1989 and 2005

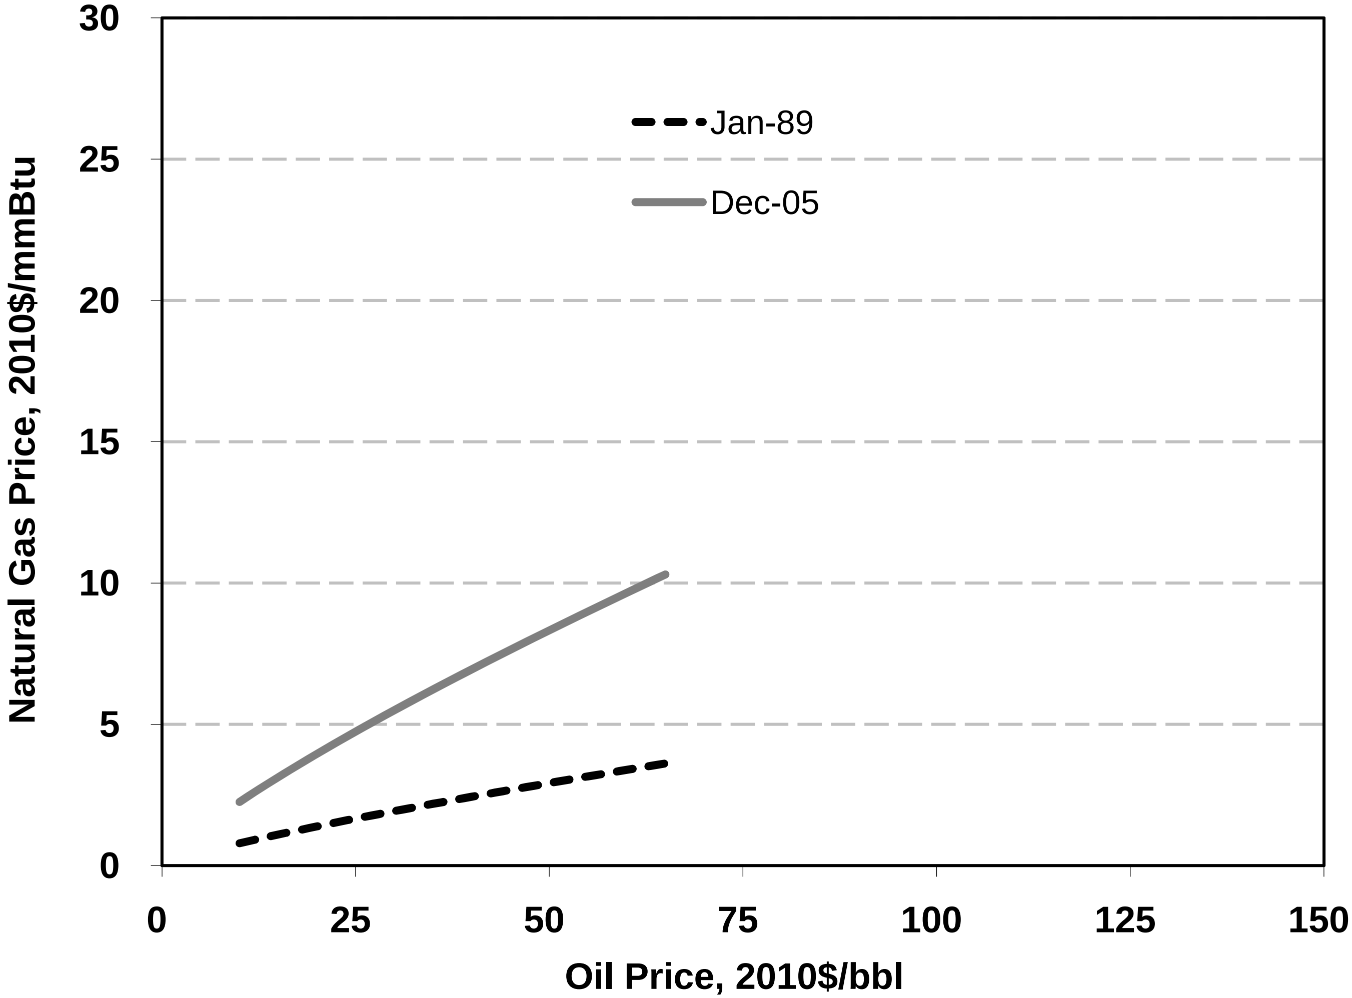

The figure graphs the estimated cointegrating relationship from Villar and Joutz (2006) at the start of their period, in January 1989, and at the end of their period, in December 2005. We only graph the relationship in the range of oil prices relevant for that window of time. However, we preserve the same scale for the overall graph as in Figure 2 in order to keep them comparable.

Figure 9. Shifting Cointegrating Relationship, 1997 -- 2009

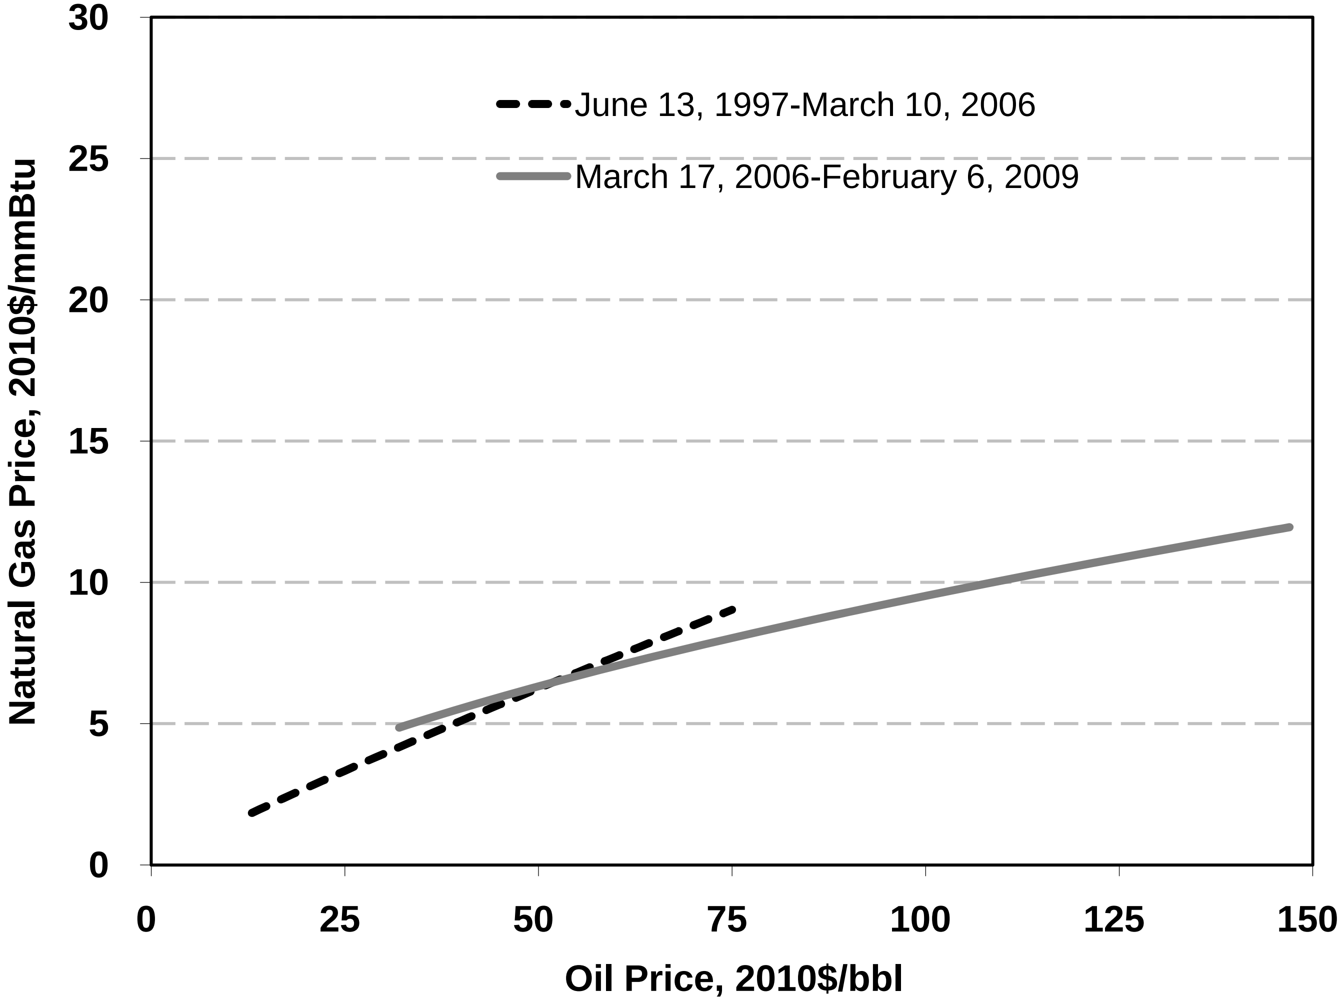

The figure graphs our estimated cointegrating relationship on two sub-periods of our sample suggested by the Gregory Hansen (1996) test. Each relationship is estimated as a part of the full Vector Error Correction Model. A Chow test confirms that the intercept and slope coefficients in the two sub-periods are different from one another.

Appendix

A1. THE DATA

Our weekly price series for natural gas and for crude oil run from January 25, 1991 through December 31, 2010. However, we only have a complete set of data for our exogenous variables—Heating and Cooling Degree Days, Deviations from Average Heating and Cooling Degree Days, Natural Gas Storage, and Shut-in Natural Gas Production—beginning on June 13, 1997. Therefore, we can discuss properties of the natural gas and oil price series beginning in 1991, but can only benefit from the information contained in the exogenous variables beginning in 1997. The primary body of statistical results reported is based on analysis using all of the exogenous variables and is therefore conducted using data from June 13, 1997 through December 31, 2010. To analyze how the fundamental relationship may have changed over time, we broke this time period up into three pieces, with one running from June 13, 1997 through April 4, 2003, the second running from April 11, 2003 through February 6, 2009, and the third running from February 13, 2009 through December 31, 2010. We will discuss the rationale behind the choice of dates in greater detail below.

The Natural Gas and the Oil Price Series

Weekly spot pricing data for West Texas Intermediate (WTI) crude oil and Henry Hub natural gas was retrieved from the Bloomberg data terminal at the MIT Sloan School of Management. Each series represents the nominal price for next-day delivery on the trade date. For the natural gas spot price series, the day-ahead price is a volume-weighted average of all trades in the Bloomberg sample on the given trade date for delivery at

Henry Hub, Louisiana. For the WTI crude oil spot price series, the day-ahead price is an arithmetic average of all sampled trades for delivery of WTI crude oil at Cushing, Oklahoma. We converted nominal dollars to real 2010 dollars as follows: we downloaded the quarterly GDP deflator series indexed to 2005, and reset the ratio so that 2010 was the benchmark. Then we linearly interpolated weekly deflator coefficients from the resulting quarterly coefficients. Finally, we divided the nominal price series by the weekly GDP deflator coefficients. All of our analyses were conducted using the *natural logs* of the two real price series. Any reference to the natural gas or oil price in this Appendix pertains to the natural logs of these real price series. We confirmed the familiar fact that each of these series appears to be non-stationary, while the first difference of each series is stationary—i.e., each series appears to be integrated of order one.[1]

[1] This was done by checking for a unit root, both with the Augmented Dickey-Fuller (ADF) test and with the Phillips-Perron test. On our statistical platform, Stata 8 Intercooled, the Augmented Dickey-Fuller test is called the Interpolated Dickey-Fuller test. The command is "dfuller". For the Phillips-Perron test, the command is "pperron."
The results are presented in Tables A1 and Table A2. Both logged real WTI crude oil and Henry Hub natural gas prices fail to reject the null hypothesis that there is a unit root and that the data are non-stationary at the 1% level of significance. The Augmented Dickey-Fuller coefficient for the natural log of the real Henry Hub price series was -3.115, with a p-value of 2.5%. The ADF coefficient for the natural log of the real WTI crude oil price was -1.409, which corresponds to a p-value of 57.8%. Identical conclusions can be drawn from the Phillips-Perron test results. For the logged real Henry Hub natural gas prices, the Z(rho) statistic is -14.448 and the Z(t) statistic is -2.780. These correspond to a p-value of 6.1%. In the case of the logged real WTI crude oil prices, the Z(rho) statistic is -3.530 and the Z(t) statistic is -1.284. These correspond to a p-value of 63.7%. The p-value is the probability that the prices observed are the prices one would expect to observe if the null hypothesis were true. Normally, to reject a null hypothesis would require a p-value below the 5% level for both tests. The ADF and Phillips-Perron test statistics thus indicate that both the Henry Hub natural gas prices and the WTI crude oil prices exhibit unit root processes.
Often a time series of data that is integrated of order one can be made stationary by differencing. That is, instead of modeling with weekly natural gas and crude oil price levels, we examine the price changes each week. The Augmented Dickey-Fuller and Phillips-Perron tests confirmed that the differenced real WTI crude oil and Henry Hub natural gas prices were stationary at the 1% significance level. For the Henry Hub natural gas prices, the first differences yielded an ADF test statistic of -28.128 and a Phillips-Perron Z(rho) statistic of -664.976. The Phillips-Perron z(t) statistic was -28.591. The equivalent statistics for the differenced real WTI crude oil price series were -28.996, -753.789, and -29.023 respectively. The results indicate that we should be able to use a regression model on the first differences of the two price series and avoid distortion due to the effect that previous price levels have on subsequent observations.

The Exogenous Variables

The models in this paper utilize a series of exogenous variables that serve as the "fundamentals" of the natural gas price. Details of each of the variables are described below.

<u>Seasonality – Heating Degree Days and Cooling Degree Days</u>

Data on Heating Degree Days (HDD) and Cooling Degree Days (CDD) were gathered from the Climate Prediction Center of the National Weather Service under the National Oceanic and Atmospheric Administration (NOAA). Both are calculated as a weighted average of weather station temperature data. HDD is the average number of degrees below 65°F. The greater the HDD figure, the greater the demand for heat. CDD is the average number of degrees above 65°F. The greater the CDD figure, the greater the demand for air conditioning. The HDD series we chose is weighted by regional gas usage. The CDD series is weighted by population. Weekly figures represent a weekly accumulation of each day's HDD and CDD.

Figures A1 and A2 depict the weekly HDD and CDD data over the 1997-2010 period. The seasonal pattern is obvious. HDD levels tend to peak about three times higher than CDD levels. The peak for HDD is in the winter, with minimum values over the summer, while the converse is true of CDD figures. During our 1997-2010 data series, the HDD variable ranged from an average high of 208 in January to an average low of 2 in July and August. The average HDD figure was 87.3, with a standard deviation of 79.8. The CDD variable typically ranged from a low of about 1 in January and February to a high of about 75 in July. CDD figures were much less volatile than HDD figures. The average CDD was 25.7, with a standard deviation of 28.5.

Unseasonal Temperature Events – deviation from normal HDD or CDD

We used deviations from normal HDD and CDD figures in order to characterize deviations from the normal seasonal temperature. These are labeled HDDDEV and CDDDEV and are calculated and reported by NOAA by subtracting, respectively, the "normal" HDD and CDD figures from the actual observations in each given week. Normal HDD and CDD figures are also reported by the NOAA and represent the average of each week's HDD and CDD reading over the 1973-2000 time period.

Figures A3 and A4 graph the HDDDEV and CDDDEV variables over our data period. The graphs also show one standard deviation both above and below the mean for each.

HDDDEV shocks tended to last about 2-3 weeks on average, reflecting the typical profile of a cold snap. Over the 1997-2010 period, deviations from normal HDD were, on average, *negative*, at -5.4, with a standard deviation of 21.0. The distribution of the deviations is skewed, with a large number of small deviations and a few large ones in the winter months. For example, the average deviation in January was -16.

CDDDEV shocks tended to last 1-2 weeks before returning to normal. From 1997-2010, the average CDDDEV figure was 2.1, with a standard deviation of 6.7. Shocks tended to be highest in June, August, and September, with an average value of 5 in each month.

Shut-in Gas Production – Hurricanes

Figure A5 plots the amount of natural gas production capacity that is curtailed in the Gulf of Mexico in million cubic feet (mmcf, equivalent to 1,030.5 mmBtu). This is always a result of hurricane activity. The shut-in natural gas production in the Gulf (SHUTIN) serves as an alternative to a “hurricane dummy” variable in order to model the one-off effects of a hurricane’s impact on the Gulf of Mexico gas industry. However, the shut-in production figure could be considered superior in one aspect: *prolonged* curtailments in production, perhaps due to damage caused by the hurricane to drilling rigs or gathering infrastructure, can be tracked according to their severity better than through the use of a binary dummy variable. The Minerals Management Service (MMS) posts weekly shut-in production statistics on its Gulf of Mexico webpage under Press Releases/Reports whenever hurricane activity prompts oil and gas producers to halt production at their offshore platforms.

Gulf of Mexico natural gas production is roughly 10% of total U.S. gas production over the course of the year (according to the Energy Information Administration (EIA) of the Department of Energy). Over the 1997-2010 period, Gulf of Mexico natural gas production ranged from about 7 to 10 Bcf (7.2 to 10.3 million mmBtu) per day (MMS website). Hurricanes on the scale of the Katrina/Rita event managed to shut in over 80% of Gulf production (MMS website press releases). Hurricane disruptions to Gulf gas production have a characteristic pattern: a large spike in the week preceding hurricane impact, followed by a gradual decline in shut-in production. This reflects the initial evacuation before hurricane impact, the immediate return to production of undamaged wells after the hurricane has passed, and the gradual return to production of rigs that were damaged to different degrees by the hurricane.

The shortest such disruption in our dataset was 34 weeks in response to Hurricane Ivan in 2004. The other two hurricane impacts in the dataset were actually combinations of two hurricane impacts in each case – Hurricanes Katrina and Rita in 2005 and Hurricanes Gustav and Ike in 2008. We thus cannot reliably calculate an "average" production impact for a single hurricane from our dataset – we have only one observation.

Storage Differentials

Natural gas is stored in various locations throughout the United States, ranging from depleted oil and/or gas fields to LNG storage facilities. The EIA collects and sums the data on storage levels on a weekly basis and reports it on its website. The STORDIFF variable is the difference between a weekly gas storage level and the 5-year running average for that date, reported in billion cubic feet (Bcf, equivalent to 1.03 million mmBtu). The average storage level from 1997-2010 was 2,332 Bcf (2.4 billion mmBtu) on any given week, with a standard deviation of 744.8 Bcf (767.53 million mmBtu). The average storage differential from 1997-2010 was 157 Bcf (162.1 million mmBtu), and the average amount of time that storage remained out of sync with the normal storage level was about 39 weeks. The median duration of a storage differential was 12 weeks. Figure A6 shows a graph of this variable.

A2. MODEL TESTS

We construct a Vector Error Correction Model (VECM) that expresses a change in the current prices for natural gas and oil in terms of past price changes. In order to determine the appropriate number of lagged effects to include in the model we first fit a vector autoregression (VAR) model using the prices and exogenous variable series in levels and then conduct a series of selection order criteria tests. The VAR model (with the exogenous variables included) is as follows:

$$P_{HH,t} = a + \sum_{i=1}^{n} b_i P_{WTI,t-i} + \sum_{i=1}^{n} c_i P_{HH,t-i} + \sum_{j=1}^{6} d_j X_{j,t} + \varepsilon_t$$

The log price of Henry Hub natural gas is determined by the previous 1 to n weeks' prices of WTI crude oil in logged real dollars per barrel ($P_{WTI,t-n}$), with each week's effect denoted by the corresponding coefficient b_i; by the previous 1 to n weeks' prices of Henry Hub natural gas in logged real dollars per mmBtu ($P_{HH,t-n}$), with each week's effect denoted by the corresponding coefficient c_i; and by the contemporaneous set of six exogenous variables (heating degree days (HDD), cooling degree days (CDD), deviations from normal HDD (HDDDEV), deviations from normal CDD (CDDDEV), shut-in natural gas production in the Gulf of Mexico (SHUTIN), and differences from average natural gas storage levels (STORDIFF). The effect of each of the six exogenous variables is denoted by the coefficient d_j. ε_t corresponds to a random error term with an expected value of zero.

The point of running the VAR selection order criteria test is to determine the number of lags, n, of previous price changes to include in the model. It involves the estimation of a series of VAR models with varying lag lengths. Each model is compared with the aim of finding the model that best explains the data for the number of parameters it uses. The tests involved penalize the use of more parameters than necessary to adequately fit the model to the data. We ran a VAR selection order criteria test with a maximum of 12 lags on the series of weekly logged real Henry Hub natural gas and WTI crude oil prices and included the exogenous variables in the VAR.[2] Using up to 12 lags in the model allows for the effects of approximately one season's duration to feed into the determination of the actual week's natural gas price. The selection order criteria tests are the Likelihood Ratio (LR) test, Akaike's Final Prediction Error (FPE), Akaike's Information Criterion (AIC) test, the Schwartz Bayesian Information Criteria (SBIC) test and the Hannan-Quinn Information Criteria (HQIC) test. Each test represents a slightly different mathematical method of statistically determining the model that achieves the best combination of fit and economy of parameters. Table A3 details the results of the VAR Selection Order Criteria tests. The Likelihood Ratio Test, the Final Prediction Error, and the Akaike Information Criteria tests all showed the closest fit at ten lags. The Hannan-Quinn Information Criteria test selected four lags, while the Schwartz Bayesian Information Criteria selected just one lag. We chose to use ten lags. A ten-lag VAR is equivalent to a nine-lag VECM.

The next step was to determine whether there is in fact a linear combination of oil and gas prices such that the series becomes stationary. This phenomenon, discussed above, is cointegration. We tested for cointegration between the oil and gas price series using the Johansen test.[3] The results of the Johansen tests are detailed in Table A4. The Johansen test indicated a rank of one (at significance of 5%) *and* a rank of zero (at 1% significance) based on the trace statistics, while the SBIC implied a rank of zero and the HQIC implied a rank of one. This is mixed evidence of a single cointegrating relationship

[2] Using the "varsoc" command on Stata 8 Intercooled.
[3] The "vecrank" command in Stata. Note that since there are only two data series here, the largest number of cointegrating relationships that can be found is one. The test is still useful in that it will identify whether the oil and gas prices are indeed cointegrated.

between Henry Hub natural gas and WTI crude oil prices if exogenous variables acting on the Henry Hub price are included.

These findings contrast with our original work using data from June 13, 1997 to February 20, 2009, in which all tests were in solid agreement of a rank of one. This provoked suspicion that the cointegrating relationship could be shifting, or even weakening, over time. For this reason, we implemented the breakpoint tests and segmented regressions discussed below.

A3. THE MODELS

The VECM

We then estimated our VECM.[4] The VECM is very similar to a VAR model, except that it includes an error-correction term and a characterization of the cointegrating relationship between two time series. Our theory is that one commodity will “lead” the other in the cointegrating relationship. The real price of the dependent commodity will commonly stray from its long-run relationship with the independent commodity in commodity markets. The error-correction term measures the “speed” at which the real price of the dependent commodity “corrects” its deviation (“error”) from this long-term equilibrium with the independent commodity by returning toward the long-run predicted relationship.

To account for natural gas fundamentals that do not affect oil prices, as well as for the volatility in natural gas that is not observed in oil prices, the VECM also incorporates

[4] In Stata, the command for a Vector Error Correction Model is “vec”.

the previously-discussed exogenous variables that act solely on our hypothesized dependent variable of real gas prices (HDD, CDD, HDDDEV, CDDDEV, SHUTIN, and STORDIFF). These fundamentals partially account for movements of the real Henry Hub natural gas price either closer to or further from the calculated long-run target.

As stated above, the difference between the actual real Henry Hub price and the long-run equilibrium real price predicted by the VECM is the "error". (This "error" is denoted by μ_t and is *not* the same as the epsilon error (ε_t), which is the difference between the actual real Henry Hub price *change* and the model's predicted real Henry Hub price *change*.) An error-correction mechanism moves the real natural gas price closer to its long-term equilibrium relationship with the real WTI price to "correct" the error by a certain coefficient each week, thus narrowing the gap between actual prices and the VECM's predicted equilibrium prices.

The mathematical representation of the VECM is as follows:

$$P_{HH,t} = \gamma + \beta P_{WTI,t} + \mu_t$$

$$\Delta P_{HH,t} = a_{HH} + \alpha_{HH}(\mu_{t-1}) + \sum_{i=1}^{n} b_{HH,i}\Delta P_{WTI,t-i} + \sum_{i=1}^{n} c_{HH,i}\Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{HH,j}X_{j,t} + \varepsilon_{HH,t}$$

$$\Delta P_{WTI,t} = a_{WTI} + \alpha_{WTI}(\mu_{t-1}) + \sum_{i=1}^{n} b_{WTI,i}\Delta P_{WTI,t-i} + \sum_{i=1}^{n} c_{WTI,i}\Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{WTI,j}X_{j,t} + \varepsilon_{WTI,t}$$

The top equation details the long-term relationship between Henry Hub and WTI prices. $P_{HH,t}$ is the logged real Henry Hub natural gas price in week t, $P_{WTI,t}$ is the logged real West Texas Intermediate crude oil price in week t, γ is a constant to be estimated, and β is a parameter to be estimated. μ_t is an error term in week t. The equation designates the "target" toward which Henry Hub prices will move over time.

The second and third equations incorporate the error correction mechanism as well as the lagged effects of the two price series on Henry Hub and WTI logged real prices, followed by the effects of the exogenous (seasonal) variables on both price series. μ_{t-1} is the lagged set of equilibrium errors in the estimated cointegrating equation. These are identical to the μ-series from the top equation, but lagged one week. X_j is the matrix of the six exogenous variables representing the fundamental drivers of the Henry Hub natural gas price. ε_t is the normal error term with a mean of zero. Finally, a, b_i, c_i, and d_j are coefficient parameters of each of the variables (each with the subscripted label to designate the commodity to which they pertain) – they will likewise be estimated by the model regression.

Note that the VECM does not simply assume that Henry Hub prices are determined by WTI prices. It mechanically returns results separately as if each price series were affected by the other. As such, the model also performs the regressions as if crude oil were the dependent variable and natural gas were the independent variable. That allows the effects of these exogenous variables to be measured on WTI crude oil, as well as an examination of the possibility that it is actually oil prices that are dependent on natural gas prices. However, the statistics alone will not allow us to conclude whether natural gas prices determine oil prices or vice versa. We need to resort to theory to draw such conclusions. In this discussion, we will focus on our hypothesized relationship, in which natural gas prices are dependent on oil prices.

Table A5 presents the results for the VECM model assuming Henry Hub prices and, alternately, WTI prices as the dependent variable. For the single cointegrating relationship, the β coefficient is 0.468, and the γ coefficient is -0.0333. The β coefficient is highly statistically significant, with a p-value of 0.0009.

The equation for the long-term relationship is as follows:

$$log(P_{HH,t}) = -0.0333 + 0.468 log(P_{WTI,t})$$

The equation for changes in the Henry Hub price has an R^2 statistic of 0.1458, while the equation for changes in the WTI price has an R^2 statistic of 0.0952. This implies that in the case of Henry Hub prices, about 14.6% of the volatility in Henry Hub prices can be described through the volatility of the WTI price and the values of the six exogenous variables. In the case of the WTI price-change equation, only about 9.5% of the volatility in WTI prices can be explained by volatility in the Henry Hub price and the values of the six exogenous variables. In line with our hypothesis, the model is better at explaining Henry Hub natural gas price movements than it is at explaining WTI crude oil price movements.

Examining the price-change equation in terms of the effect on the change in Henry Hub logged prices, we note that the modeled α-coefficient for reversion to the long-run predicted price relationship is -0.0468, with a p-value of 0.001. Based solely on this single coefficient, the implication here is that if all else held equal, any spike in Henry Hub prices from the long-run relationship would be "corrected" (be diminished) by half of the original error value in about 14.8 weeks. We call this the half-life of the error correction mechanism.

The significance of the exogenous variables representing the fundamentals of natural gas pricing as effects on the Henry Hub price varies widely. Only the variables accounting for unseasonal temperature events – HDDDEV and CDDDEV – showed p-values below the 1% threshold. Of the rest of the variables, only CDD had a p-value below 10%, which is not individually significant enough for inclusion in the model. However, the joint statistical significance of the set of six exogenous variables was high, with a Chi^2 statistic of 50.77. That corresponds to a joint p-value of 0.0000. In the end we included all of these variables in the overall Henry Hub price change equation.

We also retained the effects of the lagged price changes in Henry Hub and WTI prices on the contemporary price change in Henry Hub prices. Individually, the Henry Hub price change from one, three, five and nine weeks prior were statistically significant within the 5% p-value range, and the seventh lag was significant at the 10% level. Jointly the nine lagged price changes have a p-value of 0.0000, which is well within the 1% significance range.

Of the lagged WTI price changes, only the seventh lag was statistically significant, with a p-value of 0.007. The joint significance test of the WTI lagged price changes returned a p-value of 0.1383, which implies that the probability that the combined effects of the nine lagged WTI prices are actually zero are around 14%. However, when included with the entire set of all variables in the regression, the Chi^2 statistic is 98.91, corresponding to a statistically robust p-value of 0.0000. We thus included all of the coefficients in our modeling exercise for Henry Hub price changes.

Thus far we have focused on the second equation in the VECM: the effects of the lagged HH and WTI price changes and the six exogenous variables on the change in Henry Hub prices. However, the third equation of the VECM also provides an examination of the effects of those same variables on the change in prices for WTI crude oil. Our working hypothesis has been that the WTI crude oil price series is exogenous to the system. The VECM, however, treats WTI as a jointly endogenous element in the system. The reason we used the VECM was not to determine whether one variable or the other was exogenous to the system. Our assumptions based on observable facts lead us to believe that WTI crude oil prices are indeed free of the influence of Henry Hub natural gas prices. The VECM was run so that we could find the cointegrating (or long-run) price relationship between Henry Hub prices and WTI prices, as well as the error-correction mechanism coefficient that measures the rate at which Henry Hub prices return to the long-run relationship after deviating. Nonetheless, the statistics, while not definitive, support this assumption.

The existence of counterintuitive and spurious coefficients in the VECM model when applied to changes in the WTI crude oil price reinforce our theoretical assumption that the model does not hold any ability to estimate the price movements in WTI crude oil based on the values of the exogenous and lagged natural gas price variables. We continue to assume that WTI crude oil prices are *at least weakly* exogenous to the system. Since we do not need the added complication of assuming that WTI prices are in part determined by Henry Hub natural gas prices, we construct a modified version of our error correction model called the conditional error correction model (conditional ECM).

The conditional ECM

The conditional ECM is a VAR model in which we can measure the relationship between Henry Hub natural gas prices conditioned on the assumption that the WTI crude oil price can be taken as predetermined. In our conditional ECM, we assume that the contemporaneous change in the logged real WTI price imparts an immediate effect on the logged real price of Henry Hub natural gas. This is really an assumption that there are market factors that affect both commodities contemporaneously such that knowing the WTI price movement can allow one to infer what the Henry Hub gas prices ought to be. Furthermore, the conditional ECM uses the cointegrating error term from the VECM as an exogenous variable. The cointegrating error term is the *actual* logged real Henry Hub price minus the long-run predicted logged real Henry Hub price as predicted by the VECM for each week. Finally, we retain the six exogenous variables utilized in the original VECM.

The functional form for our conditional ECM is a vector autoregression. The error-correction term is provided solely because we took the long-run predicted cointegrating relationship estimated in the VECM as a given in the conditional ECM. Thus, the cointegrating equation is not actually estimated in our conditional ECM. The equation we estimated for changes in the Henry Hub gas price is as follows:

$$\Delta P_{HH,t} = a_{ECM} + \alpha_{ECM}(\mu_{t-1}) + b_{ECM}\Delta P_{WTI,t} + \sum_{i=1}^{n} c_{ECM,i}\Delta P_{HH,t-i} + \sum_{j=1}^{6} d_{ECM,j}X_{j,t} + \varepsilon_{ECM,t}$$

Here, in the conditional ECM, a is a constant, α is the error-correction coefficient on the lagged errors in the cointegrating relationship from the original VECM (μ_{t-1}), and b is the coefficient on the contemporaneous change in the logged WTI crude oil price. The c_i series of coefficients represent the effects of the lagged changes in the logged Henry Hub price. Since we ended up with 9 lags in the VECM, we do so again in the conditional ECM, so $n = 9$. d_j represent the coefficients for each of the exogenous variables used in the VECM (here denoted as X_j to represent HDD, CDD, HDDDEV, CDDDEV, STORDIFF, and SHUTIN). ε_t is, as before, a standard white noise error term with an expected value of zero.

Table A6 reports the results of the conditional ECM. The R^2 statistic has increased slightly to 0.1479, meaning that nearly 15% of the volatility in the Henry Hub natural gas price can be described through the fluctuations of the exogenous variables, the change in the WTI price, and the effects of the nine lagged weekly price changes in Henry Hub prices. The Chi^2 statistic has increased to 121.2 from its VECM value of 114.7.

The α coefficient for the error correction has intensified to -0.0497 and has remained statistically significant, with a p-value of 0.000. The half-life for the error correction mechanism to eliminate differences from the long-run predicted relationship is 14 weeks if the effects of the lagged Henry Hub price changes are ignored. Since the model also includes lagged price effects for Henry Hub prices, these additional coefficients also affect the rate at which the Henry Hub price can correct back to the predicted long-run relationship. When lagged effects are included, one must consider specific scenarios in which the Henry Hub price can diverge from the long-run equilibrium. This is why the half-life for the error correction differs between the scenarios depicted in Figures 16 and 17.

The statistical significance of the exogenous variables has increased slightly, with the CDD p-value improving to within the 5% significance range at 0.037. Furthermore, the signs for each are what one would expect from a change in any of the variables: increases in HDD, HDDDEV, CDDDEV, and SHUTIN provoke increases in the Henry Hub price, while increases in CDD and STORDIFF provoke decreases in the Henry Hub price. The change in the WTI price is both robust, with a coefficient of 0.2458, and also statistically significant at 0.000.

Furthermore, the Chi^2 statistics for each of the joint variable significance tests has improved considerably in every case. In the conditional ECM, the joint lagged Henry Hub price changes, the joint exogenous variables, and the combinations of lagged, WTI, and exogenous variables in joint significance tests all return p-values implying statistical significance at the 1% level or better.

A4. INSIGHTS ON THE OIL-GAS PRICE RELATIONSHIP FROM THE CONDITIONAL ECM

Comparisons to Other Models

In order to compare how this model measures up to other models of price changes, we provide Table A7, partially sourced from the IEA's *World Energy Outlook 2009* (OECD, 2009). The table compares four models in two scenarios. The four models are the Villar-Joutz model used in their 2006 paper (Villar and Joutz, 2006), the Brown-Yücel model from their 2008 paper (Brown and Yücel, 2008), and the VECM and conditional ECM as described in this paper. The two scenarios are: one in which the WTI price spikes up by 20% and holds steady thereafter (identical to that explored in Figure 16), and one in which the WTI price spikes for one period up by 20% and then returns to its original value. The table shows how the price of Henry Hub natural gas should change, percentage-wise, from its original value given the price movements in WTI. The simulated responses in natural gas prices from the VECM and the conditional ECM of this paper have about half the intensity of the responses of either Villar and Joutz or Brown and Yücel. Some possible explanations for this include the mismatch in the time periods being studied, as well as the fact that Villar and Joutz use monthly data and a time trend (as well as monthly seasonal dummy variables). While Brown and Yücel and the authors of this paper use weekly data (and HDD and CDD data as proxies for seasonal trends), there was considerable upheaval in both the crude oil and natural gas markets between June of 2007 and December of 2010, including the unprecedented explosion in crude oil prices (and subsequent collapse) and the equally dramatic decrease in natural gas prices as shale gas operations gathered momentum. It is possible that the model explored here is capturing an evolved dynamic between the two commodities that was not present during the periods of focus of earlier authors.

The Economic Significance of the Conditional ECM Coefficients

The conditional ECM provides coefficients on the effects of HDD, CDD, HDDDEV, CDDDEV, STORDIFF and SHUTIN. There are coefficients for nine weeks of lagged changes in the logged real price of Henry Hub and the contemporaneous change in the logged real price of WTI. All are used to determine how the logged real price at Henry Hub changes, but what does that mean in terms of \$/mmBtu? Table A8 details the economic effects of the exogenous variables in the conditional ECM on the change in the real price at Henry Hub in \$/mmBtu. The table provides an example of how a one-unit increase in the relevant variable will change the real price of Henry Hub natural gas from \$7/mmBtu. Note that the effect of a one-unit increase in any of the six exogenous variables is much, much smaller than the effect of a one-unit increase in either lagged real Henry Hub prices or the change in the real WTI crude oil price when prices are in natural logs. For this reason the effects of a one-standard deviation increase in each of the variables is detailed in the last column of Table A8. With few exceptions, the general magnitude of effect of a one-standard deviation increase in any of these variables on the real price of Henry Hub natural gas ranges from about 3 to 6 cents/mmBtu. The biggest outliers are the change in the WTI price, HDDDEV, CDD and CDDDEV, as well as the three-week lagged change in the real Henry Hub price. A one-standard deviation increase in HDDDEV can shift the price of Henry Hub natural gas up by 16 cents/mmBtu and the same increase in CDDDEV provokes an increase in the natural gas price of 15 cents/mmBtu when the real Henry Hub price begins at $7/mmBtu. A one-standard deviation increase in WTI prices under the same initial conditions would provoke a 10-cent/mmBtu increase in the real Henry Hub price, while a one-standard deviation increase in either CDD or the third lagged Henry Hub real price change provokes a 10-cent/mmBtu decrease in the real Henry Hub price.

Methodology for the Simulation and Seasonality Exercises

The methodology for representing the predictable seasonal price pattern due to movements in HDD and CDD (as in Figure 7) is as follows: we first averaged the HDD and CDD values over the dataset by week. We were left with a single year of weekly observations, in which each observation was the average value in that week for HDD or CDD over the 1997-2010 period. Next, we assumed a stable WTI price, and HDDDEV, CDDDEV, STORDIFF, and SHUTIN values of zero. We then used the conditional ECM coefficients for HDD and CDD and multiplied them by our average HDD and CDD values over the course of each year. We extended the simulation out for 90 years in order to settle the perturbations caused by the lagged changes in Henry Hub prices. The resulting series of predicted logged Henry Hub prices was examined. We subtracted the logged price for the average value in the series from each of the 52 weeks in the annual dataset. Then we took the natural exponent of the result. This provided us with normalized coefficients, in which the two points where the curve crosses the average level take on a value of 1 (100%), and each week has a value corresponding to that week's relationship to the average of the seasonal variability. The process is equivalent to first taking the natural exponent of the logged price series and then dividing each week's value by the average value of the seasonal benchmarks.

We used a similar methodology when using the conditional ECM to examine the predicted effects of the four exogenous variables relating to weather and supply shocks: HDDDEV, CDDDEV, STORDIFF and SHUTIN. As with our examination of the error-correction mechanism, when analyzing the effect of a single variable in the shock, we measured the effect of the shock by examining the difference in behavior of Henry Hub prices under their normal seasonal pattern (as defined by HDD and CDD above) and their behavior when affected by the shock. In order to smooth out the lagged price effects so that the seasonal pattern becomes consistent, in each case we implemented the shock in the 90th year of the simulation.

In order to characterize the magnitude of the unexplained volatility in the models, we mapped the simulated events to specific points in our 1997-2010 dataset where the data for the parameter we were examining exhibited identical behavior. In order to characterize how much more dramatic actual price changes were than the changes predicted by the model, we subtracted the predicted effects of the shock to the exogenous variable from the actual price series. This left us with an actual price series in which the effects of the exogenous variable under examination have been stripped. Changes in this modified actual price series dwarfed the predicted price effects due to shocks in the exogenous variables in every case. This provided a further illustration of the extent to which the volatility in the natural gas prices had failed to be captured by our modeling exercise.

A6. SEGMENTED TIME PERIOD MODELS

Our original WTI crude oil-Henry Hub natural gas modeling exercise used price data from June 13, 1997 to February 20, 2009. During that exercise, we were able to find strong evidence for a cointegrating relationship between WTI crude oil and Henry Hub natural gas prices regardless of any breakpoint we chose. Only the estimated parameters on each of the regressors changed. When we modeled the same breakpoints in the extended dataset, which terminates on December 31, 2010, however, we discovered (1) that the evidence of cointegration for the entire period had weakened considerably, and (2) no estimated segment following a latter-half breakpoint in the data was able to reject the null hypothesis of a unit root in the residuals of the VECM cointegrating relationship. In short, segments after the breakpoint that extended to the end of our 6/13/97 – 12/31/10 dataset could not be considered cointegrated.

We hypothesized that perhaps between February of 2009 and December of 2010 the two price series diverged dramatically enough that they could not be characterized as cointegrated. Nonetheless, we remained confident that much of the series exhibited cointegration, because the segment *before* any given breakpoint always showed strong evidence of cointegration. Our hypothesis was thus that perhaps there was some breakpoint in the 6/13/97 – 12/31/10 dataset at which the two real price series shifted from cointegrated to not cointegrated. In order to find this breakpoint, we implemented a series of tests from Gregory and Hansen (1996).

The Gregory Hansen tests are designed to test the null hypothesis of no cointegration against the alternative hypothesis of cointegration with a possible shift in the value of the estimated parameters at a single breakpoint. There is no assumed prior knowledge of where this breakpoint is supposed to occur.

They begin with the basic cointegration model with no breakpoint:

$$y_{1t} = \mu + \alpha^T y_{2t} + e_t, \qquad t = 1, \ldots . n.$$

And test it against four alternative models for breaks. The model of interest to us was the model that accounted for a shift in both the slope and the intercept of the cointegrating relationship, which they termed model #4:

$$y_{1t} = \mu_1 + \mu_2 \varphi_{t\tau} + \alpha_1^T y_{2t} + \alpha_2^T y_{2t} \varphi_{t\tau} + e_t$$

The φ coefficients are dummy variables that shift from zero to one at time $t\tau$, where t is the total number of observations in the dataset, and τ is a portion of that dataset between zero and one (only integers are used so the breakpoint can occur at a specific time step). In this case, y_{2t} is the logged real WTI crude oil price, and y_{1t} is the logged real Henry Hub natural gas price.

The Gregory-Hansen test calculates Z(α) and Z(t) Phillips-Perron style test statistics, and also an ADF test on the t-statistics from a regression of the error series on its own lags. We used the Bayesian Information Criterion to determine how many lags of the residual series were warranted for the ADF test. Details on constructing the statistics themselves are in the referenced paper.

The test is supposed to be conducted iteratively, from τ = 0.15 to τ = 0.85, but since we suspected that the cointegrating relationship deteriorated late in the June 1997 to December 2010 dataset, we extended the window for the high end of τ to 0.92. In the test, one estimates the model using OLS at each possible breakpoint and calculates the residuals. After that, the three test statistics are calculated for each breakpoint and the resulting three series of tests at each breakpoint are examined for the minimum values. These must occur in the interior of the set, and not at the endpoints. The most likely single breakpoint is the minimum statistic in the series. Gregory and Hansen then provide critical values for these tests at various significance levels. If the minimum breakpoint is smaller than the critical value they calculated at a given confidence level, the null hypothesis of no cointegration can be rejected.

For our dataset, the ADF test hit its minimum at 2/6/09 with a value of -5.349, which lies between the 5% critical value of -4.95 and the 1% critical value of -5.47. The two Phillips-Perron tests hit their minima on 5/8/09, with values that rejected the null at the 1% level in both cases. We chose 2/13/09 as the breakpoint because Gregory and Hansen tended to put the most weight on the ADF test. The two Phillips-Perron test statistics were also able to reject the null hypothesis at the 1% level on 2/13/09. This resembles the split decision for the Johansen test for cointegration on the full 1997-2010 period, in which the 5% and 1% levels disagreed as to rank. That means that we could plausibly reject the null hypothesis of no cointegration in favor of cointegration at the 2/13/09 breakpoint. This choice also worked better than when we tested the 2/6/09 breakpoint for sensitivity in the models, and we lost cointegration for the first segment if we set the breakpoint any later than 2/13/09.

We thus assumed that the correct segments are: 6/13/97 – 2/6/09 and 2/13/09 – 12/31/10. When we ran these segments, we found that the first segment was indeed cointegrated, while the second was not.[5] We will first examine the cointegrated segment.

The Cointegrated Segment: June 13, 1997 through February 6, 2009 dataset:

Since we have described this process in detail, we will only briefly present the results of the tests on the segment. The results of the unit root tests are presented in Tables A13 and A14. In each, we cannot reject the null hypothesis of a unit root in either the logged real WTI crude oil or logged real Henry Hub natural gas prices, but can do so in the differenced series. The results of the selection order criteria tests are presented in Table A19, where we select 10 lags. Table A23 presents the results of the Johansen tests. Note that in contrast to our full-period segment, the 1997-2009 segment shows that all tests are in agreement that the rank is one: cointegration is present. Tables A27 and A31 present the results of the VECM and conditional ECM for the 6/13/97-2/6/09 segment. The cointegrating relationship is now stronger: lnhh = 0.6931*lnwti – 0.8593. Furthermore, the R^2 for the conditional ECM with the shorter period is higher, at 0.165, than it was for the full period dataset. More importantly, the error correction mechanism (α) is twice as large in magnitude as it was for the conditional ECM for the full period, at -0.984. This means that in the 1997-2009 dataset, deviations in the natural gas price from the long-run cointegrating relationship are corrected twice as fast as in the estimated relationship for our full dataset.

[5] In order to examine the segmented periods and construct VECMs and conditional ECMs for them, we followed the same procedure as that described for the June 1997 – December 2010 dataset.

While we were satisfied that the 6/13/97 – 2/6/09 dataset contained strong evidence of cointegration between WTI crude oil and Henry Hub natural gas prices, we also wanted to test the hypothesis that the relationship shifted over time. To do so, we split this dataset into two equal halves, one from 6/13/97 – 4/4/03 and the other from 4/11/03 to 2/6/09, and we implemented the exact same process as described here in the Appendix for both the 1997-2010 dataset and the 1997-2009 dataset. We will describe these segments below.

The June 13, 1997 through April 4, 2003 dataset model:

The unit root tests for this model are detailed in Tables A9 and A10. SHUTIN was dropped because there were no events that shut in natural gas production capacity in this time period. Otherwise, there were no surprises in the findings.

The selection order criteria test is presented in Table A17. After comparison, we chose to use seven lags. Table A21 presents the results of the Johansen tests for cointegration. The tests uniformly suggest a rank of one – a cointegrating relationship. Table A25 presents the results of our 7-lag VECM for the 6/13/97 – 4/4/03 time period. The cointegrating relationship for this segment is lnhh = 0.726*lnwti – 1.2. The error correction term α is very statistically significant, and the coefficient is -0.17, which means that the error correction mechanism has a strong effect – much stronger than the effect calculated for either the 1997-2010 or the shorter 1997-2009 periods. The R^2 is higher than it is for either of the other periods as well, with over 20% of the variability in lnhh being explained by the regressors. We get a similar result when we fit the conditional ECM, detailed in Table A29, with one odd caveat: our total fit, R^2, is slightly worse than it was for the VECM, at 0.197. However, the statistical significance of most of the variables is higher, and the root mean squared error is also lower. We aim to contrast the model for this segment with the 4/11/03 – 2/6/09 segment to determine whether the differences in the cointegrating relationship are statistically significant. First we examine the latter half of the cointegrated period we discovered through the Gregory-Hansen tests.

<u>The April 11, 2003 through February 6, 2009 dataset model:</u>

It is in this time period that evidence begins to appear that the nature of individual time series is changing. For example, in Tables A11 and A12, the unit root tests reveal that the logged real Henry Hub natural gas price series is now able to reject the null hypothesis of a unit root at the 5% level of significance. This means that the price series can be considered stationary. Since the concept of cointegration requires that the two data series that are being compared exhibit unit root behavior, it is not clear that the VECM and conditional ECM are still appropriate models for characterizing the relationship between WTI crude oil and Henry Hub natural gas prices in this time period. Since there are other tests we can conduct if we press onward, we go ahead with the selection order criteria tests. The results are presented in Table A19. We choose ten lags, and conduct the Johansen tests for cointegration. Here, we discover that all of the tests are in agreement that the relationship between the two price series has a rank of one. According to the Johansen test, the series continue to be cointegrated. The results of the Johansen test are reproduced in Table A22.

The results of the VECM are presented in Table A26. The cointegrating relationship has shifted yet again, to lnhh = 0.4621*lnwti + 0.1969. Furthermore, the value of the α coefficient for the error correction term has diminished in magnitude to -0.0848, with a p-value of 0.011 – not robust enough to be considered significant at the 1% level. The R^2 is also lower than for the earlier segment, at 0.1461. Unit root tests on the residuals of the cointegrating relationship also reject the null hypothesis of a unit root, meaning that the residuals of the cointegrating relationship are stationary. The p-value for the ADF test was 0.0035, and the p-value for the Phillips-Perron test was 0.0042. This is another prerequisite for cointegration.

Table A30 presents the results of the conditional ECM. Here the R^2 has improved to 0.218, and the p-value of the (weakened) error correction coefficient strengthens beyond the 1% significance level. All of the joint significance tests show that, together, all groups of variables in the equation are significant beyond the 1% threshold.

It is interesting to note that the individual segments both have better fits than the full 6/13/97 through 2/6/09 cointegrated period model. But to reach a conclusion as to whether the segments are statistically significantly distinct from one another, we must conduct some formal testing.

Chow Breakpoint Tests on the Segmented Models of Cointegration

We also conducted a Chow breakpoint test on the two segments to ensure that they are distinct from one another (Chow, 1960). The null hypothesis of the Chow breakpoint test is that $\beta_{full} = \beta_{1H} = \beta_{2H}$, where β_{full} is the cointegrating relationship from 6/13/97 – 2/6/09 (the full period), β_{1H} is the cointegrating relationship from 6/13/97 – 4/4/03 (the first half of this period), and β_{2H} is the cointegrating relationship from 4/11/03 – 2/6/09. The test is quite simple. One constructs an F-statistic using the predicted values of (in our case) lnhh based on the β-coefficients of each of the segments. X_1 is the lnwti price from the first half, X_2 is the lnwti price from the second half, y_1 is the true lnhh price from the first half, and y_2 is the true lnhh price from the second half. The squared residuals are then summed as follows (Σ signifies summing over the entire segment, either X_1 or X_2, where X_1 has m elements, X_2 has n elements, and $m + n$ = all observations; p denotes the rank of the regressors (slope and intercept, rank =2)):

$$F_{p,m+n-2p} = \frac{\sum(X_1\beta_{1H} - X_1\beta_{full})^2 + \sum(X_2\beta_{2H} - X_2\beta_{full})^2}{\sum(y_1 - X_1\beta_{1H})^2 + \sum(y_2 - X_2\beta_{2H})^2} \cdot \frac{(m+n-2p)}{p}$$

We end up with an F statistic of $F_{2,609}$ = 163.47. This corresponds to a p-value of 0.00000. So we can reject the null hypothesis that there is no difference between these three alternative models. The cointegrating relationship indeed changes significantly over the 6/13/97 – 2/6/09 time period.

There is still one additional segmented time period that needs to be examined: the 2/13/09 through 12/31/10 time window for which we could not conclude that cointegraton existed between the logged real WTI crude oil price series and the logged real Henry Hub natural gas price series.

The February 13, 2009 through December 31, 2010 dataset model:

The first clue that the models for cointegration are not appropriate for characterizing the relationship between WTI crude oil prices and Henry Hub natural gas prices in this latter time window comes from the unit root test, detailed in Tables A15 and A16. Both price series reject the null hypothesis of a unit root in both sets of tests. If the series are indeed stationary, then a cointegrating relationship cannot be present, since by definition cointegration is a stationary linear relationship between two or more non-stationary data series. There are other tests we can conduct to determine whether the series are actually cointegrated or not, so we continue.

The selection order criteria tests are reproduced in Table A20. Eight lags appear to be the way to characterize the system in this time window. We run the Johansen test, and present the results in Table A24. Here we have mixed evidence of whether or not the two series are cointegrated. The SBIC and the HQIC tests find evidence of a rank of one, meaning cointegration, but according to the trace statistics, we can reject a rank of one for a rank of zero (meaning no cointegration) at the 5% level, and we can reject a rank of 2 for a rank of one at the 1% level. In order to settle the conflicting information, we run the VECM with 8 lags (see Table A28) and examine the results.

We now have a few pieces of evidence that the relationship is not cointegrated: the error correction mechanism α is not statistically significant (p-value = 0.465). The coefficient α itself is also not economically significant (0.0259) and it has the wrong sign. We have completely reversed the coefficient signs in the cointegrating relationship between crude oil and natural gas: lnhh = -3.0656*lnwti + 15.7667.

What is interesting is that the R^2 is very high (0.4858) despite these facts. This is because the coefficients for certain of the lagged HH and lagged WTI values, as well as for some of the exogenous variables, are both economically significant (large) and highly statistically significant. This is reflected by the high statistical significance we see for the "Joint Significance" tables on lnhh.

If we test the residuals of the cointegrating relationship (the long-run variable equation) for unit roots, we see that we cannot reject the null hypothesis of a unit root, meaning we are not cointegrated over this period. The ADF test statistic has a p-value of 0.1375, and the Phillips-Perron tests have a p-value of 0.1587. This is the only segment that we tested in this paper that failed to reject the null hypothesis of a unit root.

Although the evidence is clear that cointegration is not present, we fit a conditional ECM anyway, and report the results in Table A32. With a non-cointegrated system, the conditional ECM actually performs more poorly than the VECM in characterizing the lnhh series: the R^2 drops to 0.3478, the α coefficient becomes less economically significant (0.0109) and less statistically significant (0.762). It still has the wrong sign.

This last segment is clearly not cointegrated. But we need to ensure that this segment is truly statistically distinct from the segment that preceded it, so we conduct one additional Chow Breakpoint Test.

Chow Breakpoint Test on 4/11/03-2/6/09 vs. 2/13/09-12/31/10

When we run the Chow Breakpoint Test (discussed above) on the two segments, with the null being $\beta_{03\text{-}10} = \beta_{early} = \beta_{late}$, we are able to conclude that the two modeled relationships are distinct. Here $\beta_{03\text{-}10}$ is the cointegrating relationship from 4/11/03-12/31/10 (not fully laid out here because it is only used for the Chow breakpoint test, but the relationship is lnhh = 4.077*lnwti – 15.21, and it is not cointegrated. The p-values on the unit root tests of the residuals of the cointegrating relationship in this segment were 0.261 for the ADF test and 0.321 for the Phillips-Perron test statistics.) β_{early} is the segment from 4/11/03 to 2/6/09, discussed in detail above, and β_{late} is the 2/13/09-12/31/10 non-cointegrated period also discussed above. The F-statistic was $F_{2,404}$ = 713.31. This corresponds to a p-value of 0.00000. So we can reject the null hypothesis that there is no difference between these three alternative models, and we can also conclude from earlier analysis that the 4/11/03 to 2/6/09 segment is cointegrated, while the 2/13/09 to 12/31/10 segment is not.

Gregory-Hansen Test on the cointegrated period (6/13/97-2/6/09)

Up to this point, we have simply broken the time period we earlier identified (using a Gregory-Hansen test) as cointegrated into two equal halves to test the stability of the parameters over time. The Chow Breakpoint test on those halves showed that the paramteres indeed shifted over that time. But we did not conduct any formal testing to determine if the shifting parameters were due to a discrete regime shift at a specific date or due to a gradual drift of the cointegrating relationship over time. In order to provide some evidence one way or another, we also ran a Gregory-Hansen test on the 6/13/97 through 2/6/09 period, using the methodology described above to iterate across a large number of possible breakpoints.

On this test, the two Phillips-Perron test statistics and the ADF test statistic all reached a minimum on the same date: March 17, 2006. According to Gregory and Hansen's 1996 paper, the minimum point in the test represents the breakpoint for a parameter change that results in the most statistically significant bifurcated model. The values three statistics were -61.98 for the Z_α Phillips-Perron statistic, -322.00 for the Phillips-Perron Z_t statistic, and -5.51 for the ADF test statistic. So we run a segmented

VECM and Conditional ECM procedure on the two time periods, 6/13/97-3/10/06 and 3/17/06-2/6/09.

The June 13, 1997 through March 10, 2006 dataset model:

Here we begin to see evidence that the Gregory-Hansen breakpoint test for the cointegrated period, 6/13/97-2/6/09, provided useful insights. The two models which resulted provided more accurate representations of the actual movements of the Henry Hub natural gas price than the models based on our ad hoc equal-halves modeling heuristic.

First of all, Tables A33 and A34 show that for the 6/13/97-3/10/06 period, both lnhh and lnwti series are non-stationary, and are only able to reject the null hypothesis of a unit root when they are differenced. This makes the two series good candidates for cointegration models.

The selection order criteria tests, the results of which are detailed in Table A37, uniformly select seven lags in the viable lag range.[6] A Johansen test for cointegration on the segment at seven lags (Table A39) shows complete agreement among the tests of a rank of one (a single cointegrating relationship between the two price series). We fit a VECM on the 6/13/97-3/10/06 period, and report the results in Table A41. The cointegrating relationship is very steep, at lnhh = 0.91*lnwti – 1.72. This is much more dramatic than the cointegrating relationship for either the full cointegrated period (6/13/97-2/6/09) or the first-half model (6/13/97-4/4/03). The R^2 statistic, at 0.184, is similar to the first-half model, but the RMSE is smaller than either of the two alternative

[6] Remember that the VECM needs at least 2 lags in order to be implemented.

models covering this date range at 0.0902. The error correction term, α, is both economically robust, at -0.147, and statistically significant, with a p-value of 0.000.

When we fit the Conditional ECM (Table A43), all of the statistics for measuring fit improve even further. The R^2 increases to 0.209, and the RMSE improves to 0.088. This is better than either the model from 6/13/97-4/4/03 (0.197 and 0.093, respectively) or the model of the full cointegrated period of 6/13/97 through 2/6/09 (0.165 and 0.089, respectively). Furthermore, the α-coefficient has strengthened to -0.147 and remains as statistically significant as in the VECM. All of the exogenous variables except CDD are individually significant at at least the 10% level, and all combinations of lagged and exogenous variables are highly jointly significant, with a p-value of 0.0000. It would appear that the application of the Gregory-Hansen breakpoint identification process has allowed us to better characterize at least the first segment of a cointegrated period with a single breakpoint. We turn to the second segment of the cointegrated period identified by the Gregory-Hansen test on the time window of 6/13/97-2/6/09.

The March 17, 2006 through February 6, 2009 dataset model:

We apply the same methods as before to the second segment identified by the Gregory-Hansen breakpoint tests on the data through February 6, 2009. Tables A35 and A36 detail the results of the ADF and Phillips-Perron tests for unit roots in the variables. Both lnhh and lnwti fail to reject the null hypothesis of a unit root, which means that neither series is stationary in levels. Both are stationary in differences, as both differenced series reject the null hypothesis of a unit root.

The selection order criteria tests are run, with the log-likelihood test selecting 10 lags, the Final Prediction Error (FPE) and Akaike Information Criterion (AIC) selecting 4 lags, and the Hannan-Quinn Information Criterion selecting 3 lags. These are detailed in Table A38. We run the Johansen tests for cointegration on all lags, and find that only the 3-lag version is unequivocally cointegrated (with a finding of rank 1 across all tests). The 3-lag Johansen test results are reproduced in Table A40. We run the VECM with 3 lags on the 3/17/06 through 2/6/09 time window, and present the results in Table A42.

The cointegrating relationship has shifted to a more level slope than for the 6/13/97-3/10/06 time window, meaning that at higher WTI price levels, lower prices for Henry Hub natural gas are predicted than under the same conditions with the 6/13/97-3/10/06 model. The long-run equation is lnhh = 0.591*lnwti – 0.466. It is highly statistically significant, with a p-value of 0.0000. The R^2, at 0.251, and the RMSE, at 0.082, are both indicative of better fits to the actual data than the arbitrarily segmented model for the 4/4/03 through 2/6/09 window (with an R^2 of 0.205 and an RMSE of 0.084). The error correction term, α, is very strong, at -0.195, with a p-value of 0.0000. This means that approximately 20% of the Henry Hub natural gas price deviation from the long-run cointegrating relationship is corrected each week. This latter period, in which WTI crude oil prices soared from about $75 to nearly $150/bbl and back down again, appears even more strongly cointegrated than the former period.

We fit the Condtional ECM and report the results in Table A44. Here, the R^2 improves to 0.26 and the RMSE decreases even more, to 0.081.The error correction mechanism α has strengthened to -0.21 and has retained its statistical significance. In this latter period, however, the significance of the exogenous variables drops sharply. In fact, the single most potent variable is the lagged change in Henry Hub logged real natural gas prices from 2 weeks prior. The dynamics of the system appear to have shifted significantly between the period before March 17, 2006 and the period from March 17, 2006 through February 6, 2009. To be certain that a statistically significant shift has actually occurred, however, we also conducted Chow breakpoint tests to compare both the June 13, 1997 through March 10, 2006 period with the March 17, 2006 through February 6, 2009 period and the March 17, 2006 through February 6, 2009 period with the February 13, 2009 through December 31, 2010 period for statistically significant changes in the slope and intercept of the cointegrating relationship.

Chow Breakpoint Test on 6/13/97-3/10/06 vs. 3/17/06-2/6/09:

We run the Chow Breakpoint Test (discussed above) on the two segments, with the null being $\beta_{full} = \beta_{seg1} = \beta_{seg2}$. β_{full} is the slope and intercept of the cointegrating relationship from 6/13/97-2/6/09. β_{seg1} is the slope and intercept of the cointegrating relationship for the segment from 6/13/97 to 3/10/06, discussed in detail above, and β_{seg2} is the slope and intercept of the cointegrating relationship for the 3/17/06-2/6/09 segment also discussed above. The F-statistic with 2 degrees of freedom in the numerator and 609 degrees of freedom in the denominator was $F_{2,609} = 49.82$. The critical value for 1% significance with 2 degrees of freedom in the numerator and infinite degrees of freedom in the denominator is approximately 4.61, so our statistic corresponds to a p-value of 0.00000. So we can reject the null hypothesis that there is no difference between these three alternative models. The cointegrating relationship has shifted to a statistically significant degree.

Chow Breakpoint Test on 3/17/06-2/6/09 vs. 2/13/09-12/31/10:

We have a final set of Chow Breakpoint Tests to run, this time on the two segments marking the last segment of the cointegrated period and the post February 6, 2009 period which we had determined was not cointegrated. The null hypothesis is $\beta_{full} = \beta_{seg1} = \beta_{seg2}$. β_{full} is the slope and intercept of the cointegrating relationship from 3/17/06 through 12/31/10. We did not explicitly cover this model, since it is only used for this Chow test, but the cointegrating relationship was lnhh = 3.82*lnwti – 15.54. The p-value of the ADF test on the residuals of the cointegrating relationship over this period was 0.21, and the p-value of the Phillips-Perron unit root tests on the same residuals was 0.26, so we can conclude that the crude oil and natural gas price series were not cointegrated in this time window. β_{seg1} is the slope and intercept of the cointegrating relationship for the segment from 3/17/06 through 2/6/09, discussed in detail above, and β_{seg2} is the slope and intercept of the cointegrating relationship for the 2/13/09-12/31/10 segment also discussed above. The F-statistic with 2 degrees of freedom in the numerator and 251 degrees of freedom in the denominator was $F_{2,251} = 1{,}179.18$. The critical value for 1% significance with 2 degrees of freedom in the numerator and infinite degrees of freedom in the denominator is approximately 4.61, so our statistic corresponds to a p-value of 0.00000. So we can reject the null hypothesis that there is no difference between these three alternative models. We utilized the Gregory-Hansen tests for both determining the breakpoint at which the relationship between crude oil prices and natural gas prices ceases to be a cointegrating one, and also for determining where the relationship shifted while remaining cointegrated. It appears using the technique provided us with superior models than did the ad hoc selection of breakpoints.

Table A.1. Augmented Dickey-Fuller Tests				
Dates of analysis (weekly): 6/13/97 - 12/31/10				
Variable	**Levels**	**Significance %**	**1st Differences**	**Significance %**
lnhh	-3.115	2.5%	-28.128	0%
lnwti	-1.409	57.8%	-28.966	0%
hdd	-3.994	0.1%	NA	
hdddev	-16.586	0.0%	NA	
cdd	-3.919	0.2%	NA	
cdddev	-17.583	0.0%	NA	
stordiff	-1.997	28.8%	NA	
shutin	-6.612	0.0%	NA	

Table A2. Phillips-Perron Tests						
Dates of analysis (weekly): 6/13/97 - 12/31/10						
Variable	**Levels Z(rho)**	**Levels Z(t)**	**Significance Z(t) %**	**1st Differences Z(rho)**	**1st Differences Z(t)**	**Significance Z(t) %**
lnhh	-14.448	-2.780	6.1%	-664.976	-28.591	0%
lnwti	-3.530	-1.284	63.7%	-753.789	-29.023	0%
hdd	-43.461	-4.653	0%	NA	NA	
hdddev	-381.320	-16.389	0%	NA	NA	
cdd	-51.228	-5.073	0%	NA	NA	
cdddev	-506.818	-18.487	0%	NA	NA	
stordiff	-15.232	-2.800	5.8%	NA	NA	
shutin	-75.228	-6.330	0%	NA	NA	

Table A3. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 6/13/97 - 12/31/10

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	-550.19			0.017	1.62122	1.65658	1.71265
1	1630.84	4362.1	0.000	3.30E-05	-4.63460	-4.58915	-4.51705 *
2	1640.85	20.006	0.000	3.30E-05	-4.65185	-4.59630	-4.50818
3	1645.77	9.855	0.043	3.30E-05	-4.65452	-4.58887	-4.48472
4	1658.28	25.012	0.000	3.20E-05	-4.67896	-4.60321 *	-4.48304
5	1659.83	3.098	0.542	3.20E-05	-4.67192	-4.58606	-4.44988
6	1666.82	13.987	0.007	3.20E-05	-4.68052	-4.58456	-4.43235
7	1669.97	6.296	0.178	3.20E-05	-4.67807	-4.57202	-4.40378
8	1678.58	17.219	0.002	3.10E-05	-4.69132	-4.57516	-4.39091
9	1684.10	11.039	0.026	3.10E-05	-4.69568	-4.56943	-4.36915
10	1690.23	12.271 *	0.015	3.10E-05 *	-4.70182 *	-4.56546	-4.34916
11	1692.09	3.711	0.447	3.10E-05	-4.69566	-4.54920	-4.31688
12	1692.69	1.194	0.879	3.20E-05	-4.68588	-4.52932	-4.28098

Degrees of Freedom: 4, Number of Observations: 696

Endogenous Variables: log Henry Hub and log WTI prices (2010 dollars)

Exogenous Variables: HDD, CDD, HDDDEV, CDDDEV, STORDIFF, SHUTIN

Table A4. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 12/31/10

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	1687.806		16.175*	14.129	-4.367***	-4.567	-4.693
1	1694.870	0.020	2.046**	2.046	-4.359	-4.571***	-4.704
2	1695.893	0.003			-4.353	-4.569	-4.705

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

Number of Lags: 10, Number of Observations: 704

Exogenous Variables: HDD, CDD, HDDDEV, CDDDEV, STORDIFF, SHUTIN

Table A5. VECM Model including Exogenous Variables (6/13/97-12/31/10)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	26	0.0958	0.1458	114.7092	0.0000
D-lnwti	26	0.0567	0.0952	70.6978	0.0000

Long-Term Variables:	Values	P-Values	
β	0.4680	0.0009	**
γ	-0.0333		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0018	0.902	
cointegrating term (t-1) (α)	-0.0468	0.001	**
ΔPHH(t-1)	-0.1057	0.008	**
ΔPHH(t-2)	-0.0548	0.163	
ΔPHH(t-3)	-0.1498	0.000	**
ΔPHH(t-4)	-0.0414	0.283	
ΔPHH(t-5)	-0.1005	0.008	**
ΔPHH(t-6)	0.0498	0.193	
ΔPHH(t-7)	-0.0631	0.096	+
ΔPHH(t-8)	0.0184	0.628	
ΔPHH(t-9)	-0.0832	0.028	*
ΔPWTI(t-1)	0.0843	0.199	
ΔPWTI(t-2)	-0.0146	0.824	
ΔPWTI(t-3)	0.0647	0.327	
ΔPWTI(t-4)	-0.0726	0.270	
ΔPWTI(t-5)	-0.0320	0.627	
ΔPWTI(t-6)	0.0660	0.317	
ΔPWTI(t-7)	0.1793	0.007	**
ΔPWTI(t-8)	0.0744	0.260	
ΔPWTI(t-9)	0.0833	0.206	
HDD(t)	8.58E-05	0.311	
HDDDEV(t)	1.03E-03	0.000	**
CDD(t)	-4.68E-04	0.072	+
CDDDEV(t)	3.35E-03	0.000	**
STORAGE DIFF(t)	-1.94E-05	0.252	
SHUT IN(t)	4.55E-06	0.290	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	45.67	0.0000	**
Lagged WTI	13.57	0.1383	
Lagged HH & WTI	59.68	0.0000	**
Exogenous Vars.	50.77	0.0000	**
Exog + HH Lag	86.03	0.0000	**
Exog + WTI Lag	61.09	0.0000	**
Lagged + Exogs	98.91	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0076	0.369	
cointegrating term (t-1) (α)	0.0109	0.203	
ΔPHH(t-1)	0.0547	0.020	*
ΔPHH(t-2)	0.0076	0.743	
ΔPHH(t-3)	-0.0143	0.530	
ΔPHH(t-4)	0.0142	0.533	
ΔPHH(t-5)	0.0162	0.473	
ΔPHH(t-6)	0.0208	0.359	
ΔPHH(t-7)	-0.0045	0.842	
ΔPHH(t-8)	-0.0153	0.494	
ΔPHH(t-9)	-0.0309	0.167	
ΔPWTI(t-1)	-0.1124	0.004	**
ΔPWTI(t-2)	-0.0959	0.014	*
ΔPWTI(t-3)	0.0837	0.032	*
ΔPWTI(t-4)	-0.0322	0.409	
ΔPWTI(t-5)	0.0207	0.596	
ΔPWTI(t-6)	-0.0497	0.202	
ΔPWTI(t-7)	-0.0711	0.068	+
ΔPWTI(t-8)	0.1331	0.001	**
ΔPWTI(t-9)	0.0969	0.013	*
HDD(t)	5.46E-05	0.276	
HDDDEV(t)	-7.57E-05	0.528	
CDD(t)	1.43E-04	0.352	
CDDDEV(t)	3.92E-04	0.339	
STORAGE DIFF(t)	1.88E-05	0.060	+
SHUT IN(t)	-6.25E-06	0.014	*

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	10.38	0.3210	
Lagged WTI	49.84	0.0000	**
Lagged HH & WTI	56.29	0.0000	**
Exogenous Vars.	12.43	0.0531	+
Exog + HH Lag	24.23	0.0613	+
Exog + WTI Lag	62.07	0.0000	**
Lagged + Exogs	69.94	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels. Number of Observations: 698

Table A6. Conditional ECM including Exogenous Variables (6/13/97-12/31/10)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	18	0.0951	0.1479	121.1898	0.0000

Long-Term Variables:	Values	P-Values	
β	0.4680	0.0009	**
γ	-0.0333		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0044	0.750	
cointegrating term (t-1) (α)	-0.0497	0.000	**
ΔPHH(t-1)	-0.1040	0.007	**
ΔPHH(t-2)	-0.0463	0.216	
ΔPHH(t-3)	-0.1424	0.000	**
ΔPHH(t-4)	-0.0472	0.200	
ΔPHH(t-5)	-0.1090	0.003	**
ΔPHH(t-6)	0.0563	0.125	
ΔPHH(t-7)	-0.0393	0.279	
ΔPHH(t-8)	0.0376	0.299	
ΔPHH(t-9)	-0.0686	0.059	+
ΔPWTI	0.2458	0.000	**
HDD(t)	3.91E-05	0.630	
HDDDEV(t)	1.06E-03	0.000	**
CDD(t)	-5.29E-04	0.037	*
CDDDEV(t)	3.18E-03	0.000	**
STORAGE DIFF(t)	-1.94E-05	0.239	
SHUT IN(t)	4.40E-06	0.289	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	48.32	0.0000	**
Exogenous Vars.	50.35	0.0000	**
Exog + WTI	65.37	0.0000	**
Lagged + Exogs	89.70	0.0000	**
Lag + Exog + WTI	104.74	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels. Number of Observations: 698

Table A7. Impacts of Oil Price Changes on Gas Prices

Effect of a Permanent change in the price of crude oil

Researcher	Period (months)	Period (weeks)	Change in price of WTI (%)	Change in gas price at Henry Hub (%)
Villar-Joutz (1989-2005 monthly)	0	NA	20.0	5.0
	1	NA	0.0	7.8
	2	NA	0.0	9.8
	12	NA	0.0	16.0
Brown-Yücel VECM (6/13/97-6/8/07 weekly)	0	0	20.0	0.0
	0	1	0.0	3.4
	1	4	0.0	4.5
	2	8	0.0	8.5
	12	52	0.0	15.8
Ramberg-Parsons VECM (6/13/97-12/31/10 weekly)	0	0	20.0	0.0
	0	1	0.0	2.0
	1	4	0.0	1.7
	2	8	0.0	7.5
	12	52	0.0	8.4
Ramberg-Parsons Conditional ECM (6/13/97-12/31/10 weekly)	0	0	20.0	4.6
	0	1	0.0	4.3
	1	4	0.0	4.0
	2	8	0.0	4.8
	12	52	0.0	7.9

Effect of a Transitory change in the price of crude oil

Researcher	Period (months)	Period (weeks)	Change in price of WTI (%)	Change in gas price at Henry Hub (%)
Villar-Joutz (1989-2005 monthly)	0	NA	20.0	5.0
	1	NA	-16.7	2.8
	2	NA	0.0	2.1
	12	NA	0.0	0.6
Brown-Yücel VECM (6/13/97-6/8/07 weekly)	0	0	20.0	0.0
	0	1	-16.7	3.4
	1	4	0.0	-1.2
	2	8	0.0	1.2
	12	52	0.0	0.2
Ramberg-Parsons VECM (6/13/97-12/31/10 weekly)	0	0	20.0	0.0
	0	1	-16.7	2.0
	1	4	0.0	-1.5
	2	8	0.0	0.8
	12	52	0.0	0.0
Ramberg-Parsons Conditional ECM (6/13/97-12/31/10 weekly)	0	0	20.0	4.6
	0	1	-16.7	-0.3
	1	4	0.0	0.1
	2	8	0.0	0.4
	12	52	0.0	0.0

Table A8. Conditional ECM Effects of Variables on Price of Henry Hub in $/mmBtu							
Variable	**Change in HH from $7/mmBtu per One-Unit Increase in Variable**	**p-value of Exogenous Coefficient**	**% Probability that Variable's Coefficient is Actually Zero**	**Maximum Value in Data Set**	**Minimum Value in Data Set**	**Standard Deviation of Values in Data Set**	**Change in HH from $7/mmBtu per Standard Deviation Increase in Variable**
HDD	$0.00027	0.630	63.00%	272	0	79.80	$0.02
HDDDEV	$0.00741	0.000	0.00%	59	-87	20.95	$0.16
CDD	-$0.00370	0.037	3.70%	100	0	28.48	-$0.10
CDDDEV	$0.02227	0.000	0.00%	29	-19	6.68	$0.15
STORDIFF	-$0.00014	0.239	23.90%	724	-660	264.54	-$0.04
SHUTIN	$0.00003	0.289	28.90%	7941	0	970.61	$0.03
ΔPHH,t-1	-$0.69158	0.007	0.70%	0.529	-0.570	0.102	-$0.07
ΔPHH,t-2	-$0.31683	0.216	21.60%	0.529	-0.570	0.102	-$0.03
ΔPHH,t-3	-$0.92889	0.000	0.00%	0.529	-0.570	0.102	-$0.10
ΔPHH,t-4	-$0.32292	0.200	20.00%	0.529	-0.570	0.102	-$0.03
ΔPHH,t-5	-$0.72277	0.003	0.30%	0.529	-0.570	0.102	-$0.08
ΔPHH,t-6	$0.40564	0.125	12.50%	0.529	-0.570	0.102	$0.04
ΔPHH,t-7	-$0.26984	0.279	27.90%	0.529	-0.570	0.102	-$0.03
ΔPHH,t-8	$0.26856	0.299	29.90%	0.529	-0.570	0.102	$0.03
ΔPHH,t-9	-$0.46386	0.059	5.90%	0.529	-0.570	0.102	-$0.05
ΔPWTI,t	$1.95047	0.000	0.00%	0.359	-0.312	0.058	$0.10

Table A9. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 6/13/97 - 4/4/03

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.217	20.0%	-16.645	0%
lnwti	-1.755	40.3%	-18.291	0%
hdd	-2.734	6.8%	NA	NA
hdddev	-9.773	0%	NA	NA
cdd	-2.79	6.0%	NA	NA
cdddev	-14.188	0%	NA	NA
stordiff	-0.327	92.2%	NA	NA
shutin	NA	NA	NA	NA

Table A10. Phillips-Perron Tests

Dates of analysis (weekly): 6/13/97 - 4/4/03

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-8.568	-2.035	27.1%	-247.428	-16.755	0%
lnwti	-5.272	-1.609	47.9%	-297.990	-18.41	0%
hdd	-17.145	-2.983	3.7%	NA	NA	NA
hdddev	-139.664	-9.628	0%	NA	NA	NA
cdd	-20.649	-3.224	1.9%	NA	NA	NA
cdddev	-280.532	-14.624	0%	NA	NA	NA
stordiff	-3.806	-1.118	70.8%	NA	NA	NA
shutin	NA	NA	NA	NA	NA	NA

Table A11. Augmented Dickey-Fuller Tests				
Dates of analysis (weekly): 4/11/03 - 2/6/09				
Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-3.079	2.8%	-18.700	0%
lnwti	-1.704	42.9%	-20.121	0%
hdd	-2.552	10.3%	NA	NA
hdddev	-11.888	0%	NA	NA
cdd	-2.298	17.3%	NA	NA
cdddev	-9.647	0%	NA	NA
stordiff	-2.776	6.2%	NA	NA
shutin	-4.648	0%	NA	NA

Table A12. Phillips-Perron Tests						
Dates of analysis (weekly): 4/11/13 - 2/6/09						
Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-19.356	-3.057	3.0%	-317.167	-18.752	0%
lnwti	-4.500	-1.639	46.3%	-364.510	-20.003	0%
hdd	-16.494	-2.792	6.0%	NA	NA	NA
hdddev	-187.777	-11.786	0%	NA	NA	NA
cdd	-19.821	-3.144	2.4%	NA	NA	NA
cdddev	-149.274	-9.784	0%	NA	NA	NA
stordiff	-9.843	-2.831	5.4%	NA	NA	NA
shutin	-37.406	-4.479	0%	NA	NA	NA

Table A13. Augmented Dickey-Fuller Tests				
Dates of analysis (weekly): 6/13/97 - 2/6/09				
Variable	**Levels**	**Significance %**	**1st Differences**	**Significance %**
lnhh	-2.689	7.6%	-24.842	0%
lnwti	-1.607	48.0%	-27.213	0%
hdd	-3.736	0.4%	NA	NA
hdddev	-15.226	0%	NA	NA
cdd	-3.605	0.6%	NA	NA
cdddev	-16.649	0%	NA	NA
stordiff	-1.759	40.1%	NA	NA
shutin	-6.132	0%	NA	NA

Table A14. Phillips-Perron Tests						
Dates of analysis (weekly): 6/13/97 - 2/6/09						
Variable	**Levels Z(rho)**	**Levels Z(t)**	**Significance Z(t) %**	**1st Differences Z(rho)**	**1st Differences Z(t)**	**Significance Z(t) %**
lnhh	-11.421	-2.497	11.6%	-552.262	-25.020	0%
lnwti	-3.946	-1.508	52.9%	-662.244	-27.242	0%
hdd	-33.654	-4.085	0.1%	NA	NA	NA
hdddev	-326.882	-15.093	0%	NA	NA	NA
cdd	-40.274	-4.494	0%	NA	NA	NA
cdddev	-429.189	-17.242	0%	NA	NA	NA
stordiff	-12.032	-2.479	12.1%	NA	NA	NA
shutin	-63.909	-5.826	0%	NA	NA	NA

Table A15. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 2/13/09 - 12/31/10

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-3.400	1.1%	-12.307	0%
lnwti	-2.874	4.9%	-9.696	0%
hdd	-1.403	58.1%	NA	NA
hdddev	-6.885	0%	NA	NA
cdd	-1.521	52.3%	NA	NA
cdddev	-5.946	0%	NA	NA
stordiff	-2.112	24.0%	NA	NA
shutin	-4.284	0.1%	NA	NA

Table A16. Phillips-Perron Tests

Dates of analysis (weekly): 2/13/09 - 12/31/10

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-19.603	-3.314	1.4%	-117.163	-12.418	0%
lnwti	-6.679	-2.929	4.2%	-90.522	-9.706	0%
hdd	-4.910	-1.576	49.6%	NA	NA	NA
hdddev	-61.312	-6.782	0%	NA	NA	NA
cdd	-5.371	-1.646	45.9%	NA	NA	NA
cdddev	-51.595	-5.901	0%	NA	NA	NA
stordiff	-8.954	-2.392	14.4%	NA	NA	NA
shutin	-14.458	-4.458	0%	NA	NA	NA

Table A17. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 6/13/97 - 4/4/03

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	132.12			0.002	-0.82274	-0.76222	-0.67164
1	697.49	1130.700	0.000	3.20E-05	-4.66775	-4.58705 *	-4.46628 *
2	700.84	6.692	0.153	3.20E-05	-4.66327	-4.56240	-4.41144
3	706.66	11.637	0.020	3.20E-05 *	-4.67573 *	-4.55468	-4.37353
4	709.34	5.376	0.251	3.20E-05	-4.66674	-4.52552	-4.31418
5	712.36	6.021	0.198	3.20E-05	-4.65996	-4.49857	-4.25703
6	714.53	4.355	0.360	3.30E-05	-4.64748	-4.46591	-4.19418
7	721.89	14.710 *	0.005	3.20E-05	-4.67046	-4.46871	-4.16679
8	723.70	3.630	0.458	3.30E-05	-4.65549	-4.43357	-4.10146
9	724.43	1.447	0.836	3.30E-05	-4.63305	-4.39095	-4.02865
10	728.00	7.158	0.128	3.30E-05	-4.63017	-4.36790	-3.97540
11	729.71	3.407	0.492	3.40E-05	-4.61444	-4.33199	-3.90931
12	731.73	4.048	0.400	3.40E-05	-4.60090	-4.29828	-3.84540

No. of Observations = 292, Degrees of Freedom = 4.

Table A18. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 4/11/03 - 2/6/09

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	0.66			3.74E-03	0.08751	0.15581	0.25827
1	754.57	1507.800	0.000	2.70E-05	-4.82997	-4.74215	-4.61041 *
2	762.99	16.831	0.002	2.70E-05	-4.85892	-4.75159 *	-4.59057
3	766.33	6.684	0.154	2.70E-05	-4.85461	-4.72776	-4.53746
4	770.88	9.099	0.059	2.70E-05	-4.85821	-4.71184	-4.49228
5	772.13	2.509	0.643	2.70E-05	-4.84021	-4.67433	-4.42548
6	777.18	10.090	0.039	2.70E-05	-4.84706	-4.66166	-4.38354
7	778.98	3.611	0.461	2.70E-05	-4.83267	-4.62776	-4.32036
8	785.91	13.857	0.008	2.70E-05	-4.85187	-4.62744	-4.29078
9	788.68	5.529	0.237	2.70E-05	-4.84377	-4.59983	-4.23388
10	798.23	19.107 *	0.001	2.60E-05 *	-4.88019 *	-4.61673	-4.22151
11	799.64	2.831	0.587	2.70E-05	-4.86324	-4.58027	-4.15577
12	799.93	0.570	0.966	2.70E-05	-4.83888	-4.53639	-4.08262

No. of Observations = 305, Degrees of Freedom = 4.

Table A19. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 6/13/97 - 2/6/09

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	-294.68			9.64E-03	1.03411	1.07422	1.13711
1	1436.16	3461.700	0.000	3.00E-05	-4.75096	-4.69940 *	-4.61854 *
2	1443.52	14.718	0.005	2.90E-05	-4.76222	-4.69920	-4.60037
3	1449.39	11.737	0.019	2.90E-05	-4.76848	-4.69400	-4.57720
4	1454.89	11.002	0.027	2.90E-05	-4.77350	-4.68757	-4.55280
5	1456.06	2.329	0.676	2.90E-05	-4.76400	-4.66661	-4.51388
6	1460.05	7.983	0.092	2.90E-05	-4.76398	-4.65513	-4.48442
7	1462.58	5.060	0.281	2.90E-05	-4.75905	-4.63874	-4.45007
8	1470.14	15.128	0.004	2.90E-05	-4.77099	-4.63923	-4.43258
9	1473.84	7.408	0.116	2.90E-05	-4.77000	-4.62678	-4.40217
10	1481.78	15.875 *	0.003	2.90E-05 *	-4.78319 *	-4.62851	-4.38593
11	1483.95	4.341	0.362	2.90E-05	-4.77706	-4.61092	-4.35037
12	1484.98	2.052	0.726	2.90E-05	-4.76710	-4.58950	-4.31098

No. of Observations = 597, Degrees of Freedom = 4.

Table A20. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 2/13/09 - 12/31/10

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	139.05			2.74E-04	-2.52634	-2.37785	-2.15935
1	241.99	205.880	0.000	3.70E-05	-4.52511	-4.33420 *	-4.05327 *
2	244.42	4.860	0.302	3.80E-05	-4.49339	-4.26006	-3.91670
3	249.14	9.426	0.051	3.80E-05	-4.50780	-4.23204	-3.82625
4	254.81	11.355	0.023	3.70E-05	-4.54168	-4.22351	-3.75528
5	257.80	5.963	0.202	3.80E-05	-4.52111	-4.16051	-3.62986
6	262.78	9.978	0.041	3.70E-05	-4.54108	-4.13806	-3.54498
7	268.49	11.421	0.022	3.60E-05	-4.57564	-4.13019	-3.47468
8	273.68	10.377 *	0.035	3.50E-05 *	-4.59965 *	-4.11178	-3.39384
9	277.10	6.843	0.144	3.60E-05	-4.58796	-4.05767	-3.27730
10	278.81	3.419	0.490	3.80E-05	-4.54169	-3.96896	-3.12617
11	281.77	5.917	0.205	3.90E-05	-4.52064	-3.90550	-3.00027
12	283.28	3.010	0.556	4.10E-05	-4.47024	-3.81267	-2.84501

No. of Observations = 99, Degrees of Freedom = 4.

Table A21. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 4/4/03

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	723.697		28.003	27.575	-4.145	-4.428	-4.617
1	737.485	0.089	0.4276**	0.428	-4.180***	-4.486***	-4.690
2	737.698	0.001			-4.162	-4.476	-4.685

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 297, lags = 7.

Table A22. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 4/11/03 to 2/6/09

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	789.213		18.030**	14.644	-4.237***	-4.603	-4.847
1	796.535	0.047	3.386*	3.386	-4.229	-4.617***	-4.876
2	798.229	0.011			-4.222	-4.617	-4.880

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 354, lags = 10.

Table A23. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 2/6/09

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	1470.864		33.711	31.824	-4.377	-4.601	-4.744
1	1486.775	0.052	1.887**	1.887	-4.398***	-4.636***	-4.787
2	1487.719	0.003			-4.391	-4.633	-4.787

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 599, lags = 10.

Table A24. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 2/13/09 to 12/31/10

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	265.343		16.680**	15.454	-3.411	-4.067	-4.512
1	273.070	0.145	1.225*	1.225	-3.428***	-4.130***	-4.607
2	273.683	0.012			-3.394	-4.112	-4.600

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 99, lags = 8.

Table A25. VECM Model including Exogenous Variables (6/13/97-4/4/03)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	19	0.0937	0.2031	70.60894	0.0000
D-lnwti	19	0.0565	0.1452	47.0552	0.0004

Long-Term Variables:	Values	P-Values	
β	0.7261	0.0000	**
γ	-1.2007		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0010	0.960	
cointegrating term (t-1) (α)	-0.1742	0.000	**
ΔPHH(t-1)	0.0227	0.703	
ΔPHH(t-2)	-0.1241	0.036	*
ΔPHH(t-3)	-0.0158	0.785	
ΔPHH(t-4)	-0.0706	0.225	
ΔPHH(t-5)	0.0169	0.779	
ΔPHH(t-6)	0.0511	0.412	
ΔPWTI(t-1)	-0.0938	0.342	
ΔPWTI(t-2)	0.0349	0.723	
ΔPWTI(t-3)	-0.1590	0.117	
ΔPWTI(t-4)	-0.0599	0.558	
ΔPWTI(t-5)	-0.0924	0.363	
ΔPWTI(t-6)	0.1906	0.058	+
HDD(t)	3.25E-04	0.011	*
HDDDEV(t)	1.50E-03	0.000	**
CDD(t)	5.03E-05	0.896	
CDDDEV(t)	3.45E-03	0.001	**
STORAGE DIFF(t)	-5.99E-05	0.026	*
SHUT IN(t)	NA	NA	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	7.52	0.2752	
Lagged WTI	7.53	0.2743	
Lagged HH & WTI	16.95	0.1514	
Exogenous Vars.	44.09	0.0000	**
Exog + HH Lag	50.65	0.0000	**
Exog + WTI Lag	50.50	0.0000	**
Lagged + Exogs	58.34	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0230	0.064	+
cointegrating term (t-1) (α)	0.0079	0.706	
ΔPHH(t-1)	0.0814	0.023	*
ΔPHH(t-2)	0.0314	0.380	
ΔPHH(t-3)	-0.0354	0.310	
ΔPHH(t-4)	0.0526	0.135	
ΔPHH(t-5)	-0.0178	0.626	
ΔPHH(t-6)	0.0754	0.045	*
ΔPWTI(t-1)	-0.0577	0.332	
ΔPWTI(t-2)	-0.1731	0.004	**
ΔPWTI(t-3)	0.0788	0.198	
ΔPWTI(t-4)	-0.0949	0.123	
ΔPWTI(t-5)	0.0736	0.230	
ΔPWTI(t-6)	-0.1463	0.016	*
HDD(t)	1.17E-04	0.129	
HDDDEV(t)	-5.22E-05	0.769	
CDD(t)	3.52E-04	0.131	
CDDDEV(t)	1.24E-03	0.048	*
STORAGE DIFF(t)	1.97E-05	0.226	
SHUT IN(t)	NA	NA	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	13.19	0.0401	*
Lagged WTI	18.55	0.0050	**
Lagged HH & WTI	29.38	0.0035	**
Exogenous Vars.	13.31	0.0206	*
Lagged + Exogs	46.43	0.0001	**

Table A26. VECM Model including Exogenous Variables (4/11/03-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	26	0.0843	0.2051	71.99713	0.0000
D-lnwti	26	0.0574	0.1745	58.9894	0.0002

Long-Term Variables:	Values	P-Values	
β	0.4621	0.0000	**
γ	0.1969		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0034	0.870	
cointegrating term (t-1) (α)	-0.0848	0.011	*
ΔPHH(t-1)	-0.0745	0.245	
ΔPHH(t-2)	0.0037	0.952	
ΔPHH(t-3)	-0.0910	0.138	
ΔPHH(t-4)	-0.0550	0.352	
ΔPHH(t-5)	-0.1227	0.026	*
ΔPHH(t-6)	0.0294	0.587	
ΔPHH(t-7)	-0.1331	0.013	*
ΔPHH(t-8)	0.0616	0.252	
ΔPHH(t-9)	-0.1079	0.045	*
ΔPWTI(t-1)	0.0611	0.484	
ΔPWTI(t-2)	0.0434	0.623	
ΔPWTI(t-3)	0.1635	0.066	+
ΔPWTI(t-4)	0.0803	0.365	
ΔPWTI(t-5)	-0.0460	0.601	
ΔPWTI(t-6)	-0.0181	0.848	
ΔPWTI(t-7)	0.1150	0.227	
ΔPWTI(t-8)	0.1290	0.198	
ΔPWTI(t-9)	0.1919	0.052	+
HDD(t)	4.19E-05	0.720	
HDDDEV(t)	6.92E-04	0.016	*
CDD(t)	-6.64E-04	0.061	
CDDDEV(t)	3.86E-03	0.000	**
STORAGE DIFF(t)	-4.90E-05	0.039	*
SHUT IN(t)	4.84E-06	0.272	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	26.42	0.0017	**
Lagged WTI	12.21	0.2018	
Lagged HH & WTI	37.71	0.0042	**
Exogenous Vars.	26.83	0.0002	**
Exog + HH Lag	52.37	0.0000	**
Exog + WTI Lag	34.85	0.0026	**
Lagged + Exogs	60.33	0.0001	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0063	0.652	
cointegrating term (t-1) (α)	0.0453	0.046	*
ΔPHH(t-1)	0.0523	0.231	
ΔPHH(t-2)	0.0066	0.876	
ΔPHH(t-3)	-0.0030	0.943	
ΔPHH(t-4)	-0.0409	0.309	
ΔPHH(t-5)	0.0296	0.432	
ΔPHH(t-6)	-0.0055	0.882	
ΔPHH(t-7)	-0.0058	0.874	
ΔPHH(t-8)	-0.0002	0.995	
ΔPHH(t-9)	-0.1024	0.005	**
ΔPWTI(t-1)	-0.1856	0.002	**
ΔPWTI(t-2)	-0.0724	0.230	
ΔPWTI(t-3)	0.0998	0.100	+
ΔPWTI(t-4)	-0.0022	0.971	
ΔPWTI(t-5)	0.0263	0.660	
ΔPWTI(t-6)	0.0356	0.581	
ΔPWTI(t-7)	-0.1327	0.041	*
ΔPWTI(t-8)	0.1355	0.047	*
ΔPWTI(t-9)	0.2139	0.001	**
HDD(t)	3.79E-06	0.962	
HDDDEV(t)	-2.32E-04	0.235	
CDD(t)	5.44E-05	0.822	
CDDDEV(t)	-3.31E-04	0.615	
STORAGE DIFF(t)	1.20E-05	0.461	
SHUT IN(t)	-7.42E-06	0.013	*

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	10.70	0.2971	
Lagged WTI	38.13	0.0000	**
Lagged HH & WTI	46.48	0.0003	**
Exogenous Vars.	8.84	0.1830	
Lagged + Exogs	56.38	0.0002	**

Table A27. VECM Model including Exogenous Variables (6/13/97-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	26	0.0909	0.1461	98.06918	0.0000
D-lnwti	26	0.0574	0.1227	80.1048	0.0000

Long-Term Variables:	Values	P-Values	
β	0.6931	0.0000	**
γ	-0.8583		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0035	0.810	
cointegrating term (t-1) (α)	-0.0866	0.000	**
ΔPHH(t-1)	-0.0489	0.256	
ΔPHH(t-2)	-0.0696	0.102	
ΔPHH(t-3)	-0.0930	0.026	*
ΔPHH(t-4)	-0.0620	0.137	
ΔPHH(t-5)	-0.0801	0.052	+
ΔPHH(t-6)	0.0295	0.474	
ΔPHH(t-7)	-0.0661	0.106	
ΔPHH(t-8)	0.0674	0.099	+
ΔPHH(t-9)	-0.1027	0.012	*
ΔPWTI(t-1)	0.0348	0.601	
ΔPWTI(t-2)	0.0357	0.593	
ΔPWTI(t-3)	0.0569	0.401	
ΔPWTI(t-4)	-0.0013	0.985	
ΔPWTI(t-5)	-0.0170	0.800	
ΔPWTI(t-6)	0.1003	0.145	
ΔPWTI(t-7)	0.1158	0.094	+
ΔPWTI(t-8)	0.0372	0.601	
ΔPWTI(t-9)	0.0990	0.162	
HDD(t)	1.20E-04	0.166	
HDDDEV(t)	1.01E-03	0.000	**
CDD(t)	-4.09E-04	0.126	
CDDDEV(t)	3.39E-03	0.000	**
STORAGE DIFF(t)	-3.58E-05	0.036	*
SHUT IN(t)	5.04E-06	0.208	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	24.72	0.0033	**
Lagged WTI	7.67	0.5677	
Lagged HH & WTI	31.63	0.0243	*
Exogenous Vars.	52.90	0.0000	**
Exog + HH Lag	74.78	0.0000	**
Exog + WTI Lag	57.01	0.0000	**
Lagged + Exogs	79.59	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0093	0.308	
cointegrating term (t-1) (α)	0.0324	0.010	**
ΔPHH(t-1)	0.0739	0.006	**
ΔPHH(t-2)	0.0210	0.435	
ΔPHH(t-3)	-0.0261	0.322	
ΔPHH(t-4)	0.0022	0.932	
ΔPHH(t-5)	0.0140	0.592	
ΔPHH(t-6)	0.0338	0.193	
ΔPHH(t-7)	0.0128	0.621	
ΔPHH(t-8)	0.0057	0.824	
ΔPHH(t-9)	-0.0314	0.226	
ΔPWTI(t-1)	-0.1345	0.001	**
ΔPWTI(t-2)	-0.1020	0.016	*
ΔPWTI(t-3)	0.0976	0.022	*
ΔPWTI(t-4)	-0.0357	0.402	
ΔPWTI(t-5)	0.0309	0.466	
ΔPWTI(t-6)	-0.0454	0.297	
ΔPWTI(t-7)	-0.1062	0.015	*
ΔPWTI(t-8)	0.0935	0.037	*
ΔPWTI(t-9)	0.1350	0.003	**
HDD(t)	4.76E-05	0.385	
HDDDEV(t)	-1.73E-04	0.177	
CDD(t)	2.03E-04	0.230	
CDDDEV(t)	4.25E-04	0.351	
STORAGE DIFF(t)	2.62E-05	0.015	*
SHUT IN(t)	-7.43E-06	0.003	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	12.67	0.1779	
Lagged WTI	49.17	0.0000	**
Lagged HH & WTI	56.21	0.0000	**
Exogenous Vars.	19.26	0.0037	**
Lagged + Exogs	76.80	0.0000	**

Table A28. VECM Model including Exogenous Variables (2/13/09-12/31/10)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	22	0.1058	0.4858	72.73354	0.0000
D-lnwti	22	0.0452	0.3265	37.3320	0.0217

Long-Term Variables:	Values	P-Values	
β	-3.0656	0.0000	**
γ	15.7667		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0763	0.293	
cointegrating term (t-1) (α)	0.0259	0.465	
ΔPHH(t-1)	-0.3520	0.002	**
ΔPHH(t-2)	-0.1682	0.119	
ΔPHH(t-3)	-0.3848	0.000	**
ΔPHH(t-4)	-0.0543	0.605	
ΔPHH(t-5)	-0.1173	0.228	
ΔPHH(t-6)	0.1150	0.244	
ΔPHH(t-7)	0.0639	0.520	
ΔPWTI(t-1)	0.3296	0.172	
ΔPWTI(t-2)	-0.5912	0.012	*
ΔPWTI(t-3)	0.1073	0.639	
ΔPWTI(t-4)	-0.6028	0.006	**
ΔPWTI(t-5)	-0.4633	0.047	*
ΔPWTI(t-6)	-0.2014	0.312	
ΔPWTI(t-7)	0.1856	0.353	
HDD(t)	1.25E-04	0.700	
HDDDEV(t)	1.88E-03	0.032	*
CDD(t)	-1.63E-03	0.058	+
CDDDEV(t)	6.38E-03	0.002	**
STORAGE DIFF(t)	4.46E-04	0.000	**
SHUT IN(t)	-1.01E-05	0.903	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	32.53	0.0000	**
Lagged WTI	21.35	0.0033	**
Lagged HH & WTI	58.65	0.0000	**
Exogenous Vars.	24.38	0.0004	**
Exog + HH Lag	45.42	0.0000	**
Exog + WTI Lag	43.29	0.0000	**
Lagged + Exogs	72.05	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0376	0.226	
cointegrating term (t-1) (α)	-0.0526	0.001	**
ΔPHH(t-1)	0.0108	0.825	
ΔPHH(t-2)	-0.0150	0.745	
ΔPHH(t-3)	0.0197	0.661	
ΔPHH(t-4)	0.0380	0.397	
ΔPHH(t-5)	0.0389	0.350	
ΔPHH(t-6)	0.0002	0.996	
ΔPHH(t-7)	-0.0197	0.644	
ΔPWTI(t-1)	-0.0719	0.486	
ΔPWTI(t-2)	-0.0880	0.383	
ΔPWTI(t-3)	-0.1193	0.223	
ΔPWTI(t-4)	-0.1027	0.275	
ΔPWTI(t-5)	-0.1430	0.151	
ΔPWTI(t-6)	-0.2751	0.001	**
ΔPWTI(t-7)	-0.2068	0.016	*
HDD(t)	5.86E-05	0.673	
HDDDEV(t)	1.26E-04	0.737	
CDD(t)	-4.15E-05	0.910	
CDDDEV(t)	-6.52E-05	0.941	
STORAGE DIFF(t)	-4.03E-05	0.451	
SHUT IN(t)	-5.67E-05	0.109	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	2.57	0.9218	
Lagged WTI	13.76	0.0557	+
Lagged HH & WTI	18.55	0.1827	
Exogenous Vars.	4.72	0.5805	
Lagged + Exogs	26.14	0.1611	

Table A29. Conditional ECM Including Exogenous Variables (6/13/97 - 4/4/03)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	14	0.093205	0.197	72.85282	0.0000

Long-Term Variables:	Values	P-Values	
β	0.7261	0.0000	**
γ	-1.2007		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0073	0.714	
cointegrating term (t-1) (α)	-0.1732	0.000	**
ΔPHH(t-1)	-0.0003	0.996	
ΔPHH(t-2)	-0.1325	0.017	*
ΔPHH(t-3)	-0.0278	0.608	
ΔPHH(t-4)	-0.0888	0.107	
ΔPHH(t-5)	0.0039	0.946	
ΔPHH(t-6)	0.0462	0.436	
ΔPWTI	0.2269	0.015	*
HDD(t)	2.76E-04	0.024	*
HDDDEV(t)	1.50E-03	0.000	**
CDD(t)	-7.56E-05	0.839	
CDDDEV(t)	2.92E-03	0.004	**
STORAGE DIFF(t)	-6.47E-05	0.013	*
SHUT IN(t)	NA	NA	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	9.44	0.1501	
Exogenous Vars.	45.65	0.0000	**
Exog + WTI	51.63	0.0000	**
Lagged + Exogs	53.30	0.0000	**
Lag + Exog + WTI	59.99	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Table A30. Conditional ECM Including Exogenous Variables (4/11/03-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	18	0.082398	0.2182	85.10134	0.0000

Long-Term Variables:	Values	P-Values	
β	0.4621	0.0000	**
γ	0.1969		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0059	0.759	
cointegrating term (t-1) (α)	-0.0902	0.004	**
ΔPHH(t-1)	-0.0708	0.231	
ΔPHH(t-2)	0.0102	0.858	
ΔPHH(t-3)	-0.0657	0.237	
ΔPHH(t-4)	-0.0106	0.844	
ΔPHH(t-5)	-0.1309	0.010	**
ΔPHH(t-6)	0.0385	0.432	
ΔPHH(t-7)	-0.1145	0.019	*
ΔPHH(t-8)	0.1004	0.038	*
ΔPHH(t-9)	-0.0390	0.427	
ΔPWTI	0.3377	0.000	**
HDD(t)	-9.99E-06	0.926	
HDDDEV(t)	7.58E-04	0.005	**
CDD(t)	-6.69E-04	0.045	*
CDDDEV(t)	3.97E-03	0.000	**
STORAGE DIFF(t)	-3.60E-05	0.091	+
SHUT IN(t)	3.84E-06	0.332	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	28.15	0.0009	**
Exogenous Vars.	26.74	0.0001	**
Exog + WTI	42.82	0.0000	**
Lagged + Exogs	54.45	0.0000	**
Lag + Exog + WTI	70.65	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Table A31. Conditional ECM Including Exogenous Variables (6/13/97-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	18	0.089292	0.1650	118.3864	0.0000

Long-Term Variables:	Values	P-Values	
β	0.6931	0.0000	**
γ	-0.8583		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0015	0.912	
cointegrating term (t-1) (α)	-0.0984	0.000	**
ΔPHH(t-1)	-0.0545	0.181	
ΔPHH(t-2)	-0.0584	0.141	
ΔPHH(t-3)	-0.0697	0.073	+
ΔPHH(t-4)	-0.0548	0.160	
ΔPHH(t-5)	-0.0807	0.037	*
ΔPHH(t-6)	0.0362	0.350	
ΔPHH(t-7)	-0.0413	0.281	
ΔPHH(t-8)	0.0850	0.026	*
ΔPHH(t-9)	-0.0833	0.030	*
ΔPWTI	0.2866	0.000	**
HDD(t)	7.93E-05	0.336	
HDDDEV(t)	1.05E-03	0.000	**
CDD(t)	-4.60E-04	0.073	+
CDDDEV(t)	3.23E-03	0.000	**
STORAGE DIFF(t)	-3.97E-05	0.016	*
SHUT IN(t)	5.51E-06	0.149	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	24.83	0.0032	**
Exogenous Vars.	56.57	0.0000	**
Exog + WTI	74.52	0.0000	**
Lagged + Exogs	79.64	0.0000	**
Lag + Exog + WTI	98.66	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Table A32. Conditional ECM Including Exogenous Variables (2/13/09-12/31/10)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	16	0.114735	0.3478	52.79584	0.0000

Long-Term Variables:	Values	P-Values	
β	-3.0656	0.0000	**
γ	15.7667		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0961	0.166	
cointegrating term (t-1) (α)	0.0109	0.762	
ΔPHH(t-1)	-0.3112	0.003	**
ΔPHH(t-2)	-0.1075	0.289	
ΔPHH(t-3)	-0.3984	0.000	**
ΔPHH(t-4)	-0.0770	0.455	
ΔPHH(t-5)	-0.1366	0.153	
ΔPHH(t-6)	0.1051	0.273	
ΔPHH(t-7)	0.0481	0.617	
ΔPWTI	0.2050	0.400	
HDD(t)	2.50E-04	0.421	
HDDDEV(t)	1.68E-03	0.040	*
CDD(t)	-1.14E-03	0.165	
CDDDEV(t)	5.59E-03	0.005	**
STORAGE DIFF(t)	3.58E-04	0.003	**
SHUT IN(t)	-5.23E-05	0.507	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	38.53	0.0000	**
Exogenous Vars.	21.68	0.0014	**
Exog + WTI	22.96	0.0017	**
Lagged + Exogs	51.83	0.0000	**
Lag + Exog + WTI	52.11	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Table A33. Augmented Dickey-Fuller Tests				
Dates of analysis (weekly): 6/13/97 - 3/10/06				
Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.251	18.8%	-20.743	0.0%
lnwti	-1.188	67.9%	-22.566	0.0%
hdd	-3.322	1.4%	NA	NA
hdddev	-12.574	0.0%	NA	NA
cdd	-3.257	1.7%	NA	NA
cdddev	-15.348	0.0%	NA	NA
stordiff	-0.773	82.7%	NA	NA
shutin	-5.391	0.0%	NA	NA

Table A34. Phillips-Perron Tests						
Dates of analysis (weekly): 6/13/97 - 3/10/06						
Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-8.011	-2.041	26.9%	-381.169	-20.915	0.0%
lnwti	-2.651	-0.967	76.5%	-444.562	-22.772	0.0%
hdd	-25.747	-3.634	0.5%	NA	NA	NA
hdddev	-225.603	-12.394	0.0%	NA	NA	NA
cdd	-30.395	-3.913	0.2%	NA	NA	NA
cdddev	-359.311	-15.960	0.0%	NA	NA	NA
stordiff	-7.352	-1.692	43.5%	NA	NA	NA
shutin	-48.108	-5.039	0.0%	NA	NA	NA

Table A35. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 3/17/06 - 2/6/09

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.127	23.4%	-14.001	0.0%
lnwti	-0.693	84.9%	-14.736	0.0%
hdd	-1.728	41.7%	NA	NA
hdddev	-8.830	0.0%	NA	NA
cdd	-1.567	50.0%	NA	NA
cdddev	-6.817	0.0%	NA	NA
stordiff	-1.772	39.4%	NA	NA
shutin	-3.016	3.3%	NA	NA

Table A36. Phillips-Perron Tests

Dates of analysis (weekly): 3/17/06 - 2/6/09

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-12.266	-2.313	16.8%	-189.225	-13.891	0.0%
lnwti	-1.773	-0.623	86.6%	-195.648	-14.547	0.0%
hdd	-7.187	-1.780	39.0%	NA	NA	NA
hdddev	-104.828	-8.861	0.0%	NA	NA	NA
cdd	-8.922	-2.094	24.7%	NA	NA	NA
cdddev	-71.695	-6.804	0.0%	NA	NA	NA
stordiff	-4.475	-1.852	35.5%	NA	NA	NA
shutin	-18.360	-3.109	2.6%	NA	NA	NA

Table A37. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 6/13/97 - 3/10/06

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	5.26			0.004	0.03929	0.09013	0.16822
1	1107.42	2204.300	0.000	2.60E-05	-4.89628	-4.83091 *	-4.73051 *
2	1110.42	5.991	0.200	2.60E-05	-4.89176	-4.81187	-4.68916
3	1119.97	19.110	0.001	2.50E-05	-4.91673	-4.82232	-4.67729
4	1122.52	5.089	0.278	2.50E-05	-4.91019	-4.80125	-4.63392
5	1127.42	9.810	0.044	2.50E-05	-4.91426	-4.79079	-4.60115
6	1131.48	8.122	0.087	2.50E-05	-4.91453	-4.77654	-4.56458
7	1137.94	12.920 *	0.012	2.50E-05 *	-4.92559 *	-4.77307	-4.53880
8	1140.09	4.289	0.368	2.50E-05	-4.91725	-4.75021	-4.49363
9	1141.03	1.877	0.758	2.50E-05	-4.90349	-4.72192	-4.44303
10	1144.74	7.427	0.115	2.50E-05	-4.90220	-4.70611	-4.40491
11	1146.04	2.605	0.626	2.60E-05	-4.89008	-4.67946	-4.35595
12	1148.57	5.046	0.283	2.60E-05	-4.88344	-4.65830	-4.31247

No. of Observations = 445, Degrees of Freedom = 4.

Table A38. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly):3/17/06 - 2/6/09

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwartz Bayesian Information Criterion (SBIC)
0	80.85			1.42E-03	-0.87959	-0.76645	-0.60108
1	356.28	550.860	0.000	4.00E-05	-4.45102	-4.30555	-4.09293 *
2	362.83	13.098	0.011	3.90E-05	-4.48456	-4.30676	-4.04689
3	372.71	19.764	0.001	3.60E-05	-4.56195	-4.35183 *	-4.04471
4	377.12	8.813	0.066	3.60E-05 *	-4.56730 *	-4.32485	-3.97048
5	378.75	3.271	0.514	3.70E-05	-4.53619	-4.26141	-3.85979
6	382.40	7.291	0.121	3.70E-05	-4.53152	-4.22442	-3.77555
7	384.56	4.319	0.365	3.80E-05	-4.50731	-4.16788	-3.67176
8	390.79	12.474	0.014	3.70E-05	-4.53674	-4.16499	-3.62162
9	392.31	3.026	0.553	3.80E-05	-4.50402	-4.09994	-3.50932
10	400.85	17.079 *	0.002	3.60E-05	-4.56375	-4.12734	-3.48947
11	404.19	6.699	0.153	3.60E-05	-4.55519	-4.08645	-3.40134
12	406.34	4.296	0.367	3.70E-05	-4.53082	-4.02976	-3.29739

No. of Observations = 152, Degrees of Freedom = 4.

Table A39. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 3/10/06

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	1134.639		38.800	38.578	-4.527	-4.737	-4.874
1	1153.928	0.082	0.2221**	0.222	-4.572***	-4.799***	-4.946
2	1154.039	0.000			-4.559	-4.791	-4.942

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 450, lags = 7.

Table A40. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 3/17/06 to 2/6/09

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	362.320		20.777	20.771	-4.040	-4.300	-4.478
1	372.705	0.128	0.0062**	0.006	-4.078***	-4.373***	-4.575
2	372.708	0.000			-4.045	-4.352	-4.562

* = significant at the 5% level, ** = significant at the 1% level, *** = "best fit" according to various criteria

No. of Observations = 152, lags = 3.

Table A41. VECM Model including Exogenous Variables (6/13/97-3/10/06)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	20	0.0902	0.1844	97.22279	0.0000
D-lnwti	20	0.0536	0.1148	55.7409	0.0000

Long-Term Variables:	Values	P-Values	
β	0.9076	0.0000	**
γ	-1.7182		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0018	0.910	
cointegrating term (t-1) (α)	-0.1471	0.000	**
ΔPHH(t-1)	-0.0047	0.923	
ΔPHH(t-2)	-0.1398	0.003	**
ΔPHH(t-3)	-0.0201	0.667	
ΔPHH(t-4)	-0.0894	0.055	+
ΔPHH(t-5)	0.0128	0.784	
ΔPHH(t-6)	0.0007	0.988	
ΔPWTI(t-1)	-0.0873	0.288	
ΔPWTI(t-2)	0.0140	0.865	
ΔPWTI(t-3)	-0.0645	0.433	
ΔPWTI(t-4)	-0.0473	0.566	
ΔPWTI(t-5)	-0.1318	0.107	
ΔPWTI(t-6)	0.0753	0.353	
HDD(t)	2.25E-04	0.024	*
HDDDEV(t)	1.36E-03	0.000	**
CDD(t)	-2.16E-06	0.994	
CDDDEV(t)	3.08E-03	0.000	**
STORAGE DIFF(t)	-4.82E-05	0.014	*
SHUT IN(t)	1.20E-05	0.010	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	11.52	0.0735	+
Lagged WTI	4.93	0.5524	
Lagged HH & WTI	18.41	0.1038	
Exogenous Vars.	62.36	0.0000	**
Exog + HH Lag	73.67	0.0000	**
Exog + WTI Lag	65.74	0.0000	**
Lagged + Exogs	78.38	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0174	0.068	+
cointegrating term (t-1) (α)	0.0153	0.314	
ΔPHH(t-1)	0.0529	0.065	+
ΔPHH(t-2)	0.0237	0.404	
ΔPHH(t-3)	-0.0410	0.139	
ΔPHH(t-4)	0.0346	0.211	
ΔPHH(t-5)	-0.0098	0.721	
ΔPHH(t-6)	0.0578	0.036	*
ΔPWTI(t-1)	-0.0523	0.284	
ΔPWTI(t-2)	-0.1862	0.000	**
ΔPWTI(t-3)	0.0865	0.077	+
ΔPWTI(t-4)	-0.1185	0.016	*
ΔPWTI(t-5)	0.0664	0.172	
ΔPWTI(t-6)	-0.1376	0.004	**
HDD(t)	8.72E-05	0.141	
HDDDEV(t)	-8.80E-05	0.523	
CDD(t)	3.54E-04	0.050	*
CDDDEV(t)	4.65E-04	0.324	
STORAGE DIFF(t)	1.82E-05	0.117	
SHUT IN(t)	-2.38E-06	0.393	

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	12.32	0.0552	+
Lagged WTI	30.85	0.0000	**
Lagged HH & WTI	39.63	0.0001	**
Exogenous Vars.	11.79	0.0669	+
Lagged + Exogs	53.37	0.0000	**

Table A42. VECM Model including Exogenous Variables (3/17/06-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	12	0.0821	0.2514	47.0042	0.0000
D-lnwti	12	0.0673	0.165	27.6705	0.0062

Long-Term Variables:	Values	P-Values	
β	0.5906	0.0000	**
γ	-0.4663		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0039	0.892	
cointegrating term (t-1) (α)	-0.1947	0.000	**
ΔPHH(t-1)	-0.0439	0.606	
ΔPHH(t-2)	0.3246	0.000	**
ΔPWTI(t-1)	-0.0287	0.777	
ΔPWTI(t-2)	-0.0997	0.321	
HDD(t)	-2.01E-05	0.897	
HDDDEV(t)	1.94E-04	0.622	
CDD(t)	-9.81E-04	0.046	*
CDDDEV(t)	3.98E-03	0.005	**
STORAGE DIFF(t)	-9.18E-06	0.789	
SHUT IN(t)	-1.10E-05	0.098	+

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	19.80	0.0001	**
Lagged WTI	0.98	0.6113	
Lagged HH & WTI	20.15	0.0005	**
Exogenous Vars.	13.68	0.0334	*
Exog + HH Lag	36.06	0.0000	**
Exog + WTI Lag	14.39	0.0721	+
Lagged + Exogs	36.70	0.0001	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0123	0.602	
cointegrating term (t-1) (α)	0.0619	0.120	
ΔPHH(t-1)	0.0765	0.272	
ΔPHH(t-2)	0.0973	0.135	
ΔPWTI(t-1)	-0.2836	0.001	**
ΔPWTI(t-2)	-0.1241	0.132	
HDD(t)	-8.57E-05	0.500	
HDDDEV(t)	-2.20E-04	0.495	
CDD(t)	-1.25E-04	0.757	
CDDDEV(t)	2.31E-04	0.842	
STORAGE DIFF(t)	3.89E-05	0.167	
SHUT IN(t)	-1.62E-05	0.003	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	2.72	0.2562	
Lagged WTI	12.07	0.0024	**
Lagged HH & WTI	13.10	0.0108	*
Exogenous Vars.	14.14	0.0281	*
Lagged + Exogs	21.92	0.0155	*

Table A43. Conditional ECM Including Exogenous Variables (6/13/97 - 3/10/06)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	15	0.088324	0.2094	119.1844	0.0000

Long-Term Variables:	Values	P-Values	
β	0.9076	0.0000	**
γ	-1.7182		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0057	0.713	
cointegrating term (t-1) (α)	-0.1490	0.000	**
ΔPHH(t-1)	-0.0313	0.490	
ΔPHH(t-2)	-0.1414	0.001	**
ΔPHH(t-3)	-0.0172	0.693	
ΔPHH(t-4)	-0.1042	0.017	*
ΔPHH(t-5)	-0.0073	0.867	
ΔPHH(t-6)	-0.0111	0.797	
ΔPWTI	0.3352	0.000	**
HDD(t)	1.82E-04	0.056	+
HDDDEV(t)	1.37E-03	0.000	**
CDD(t)	-1.58E-04	0.586	
CDDDEV(t)	2.87E-03	0.000	**
STORAGE DIFF(t)	-5.46E-05	0.004	**
SHUT IN(t)	1.27E-05	0.005	**

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	14.37	0.0258	*
Exogenous Vars.	64.60	0.0000	**
Exog + WTI	82.71	0.0000	**
Lagged + Exogs	77.68	0.0000	**
Lag + Exog + WTI	96.14	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Table A44. Conditional ECM Including Exogenous Variables (3/17/06-2/6/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	11	0.081322	0.2603	53.49019	0.0000

Long-Term Variables:	Values	P-Values	
β	0.5906	0.0000	**
γ	-0.4663		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0014	0.958	
cointegrating term (t-1) (α)	-0.2061	0.000	**
ΔPHH(t-1)	-0.0580	0.469	
ΔPHH(t-2)	0.3036	0.000	**
ΔPWTI	0.1646	0.081	+
HDD(t)	-2.63E-06	0.986	
HDDDEV(t)	2.22E-04	0.554	
CDD(t)	-9.62E-04	0.040	*
CDDDEV(t)	3.97E-03	0.003	**
STORAGE DIFF(t)	-1.68E-05	0.609	
SHUT IN(t)	-7.67E-06	0.210	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	20.06	0.0000	**
Exogenous Vars.	12.86	0.0247	*
Exog + WTI	15.69	0.0155	*
Lagged + Exogs	35.84	0.0000	**
Lag + Exog + WTI	39.76	0.0000	**

+ = 0.1, * = 0.05, ** = 0.01 significance levels

Figure A1. US Average Weekly Heating Degree Days (HDD)

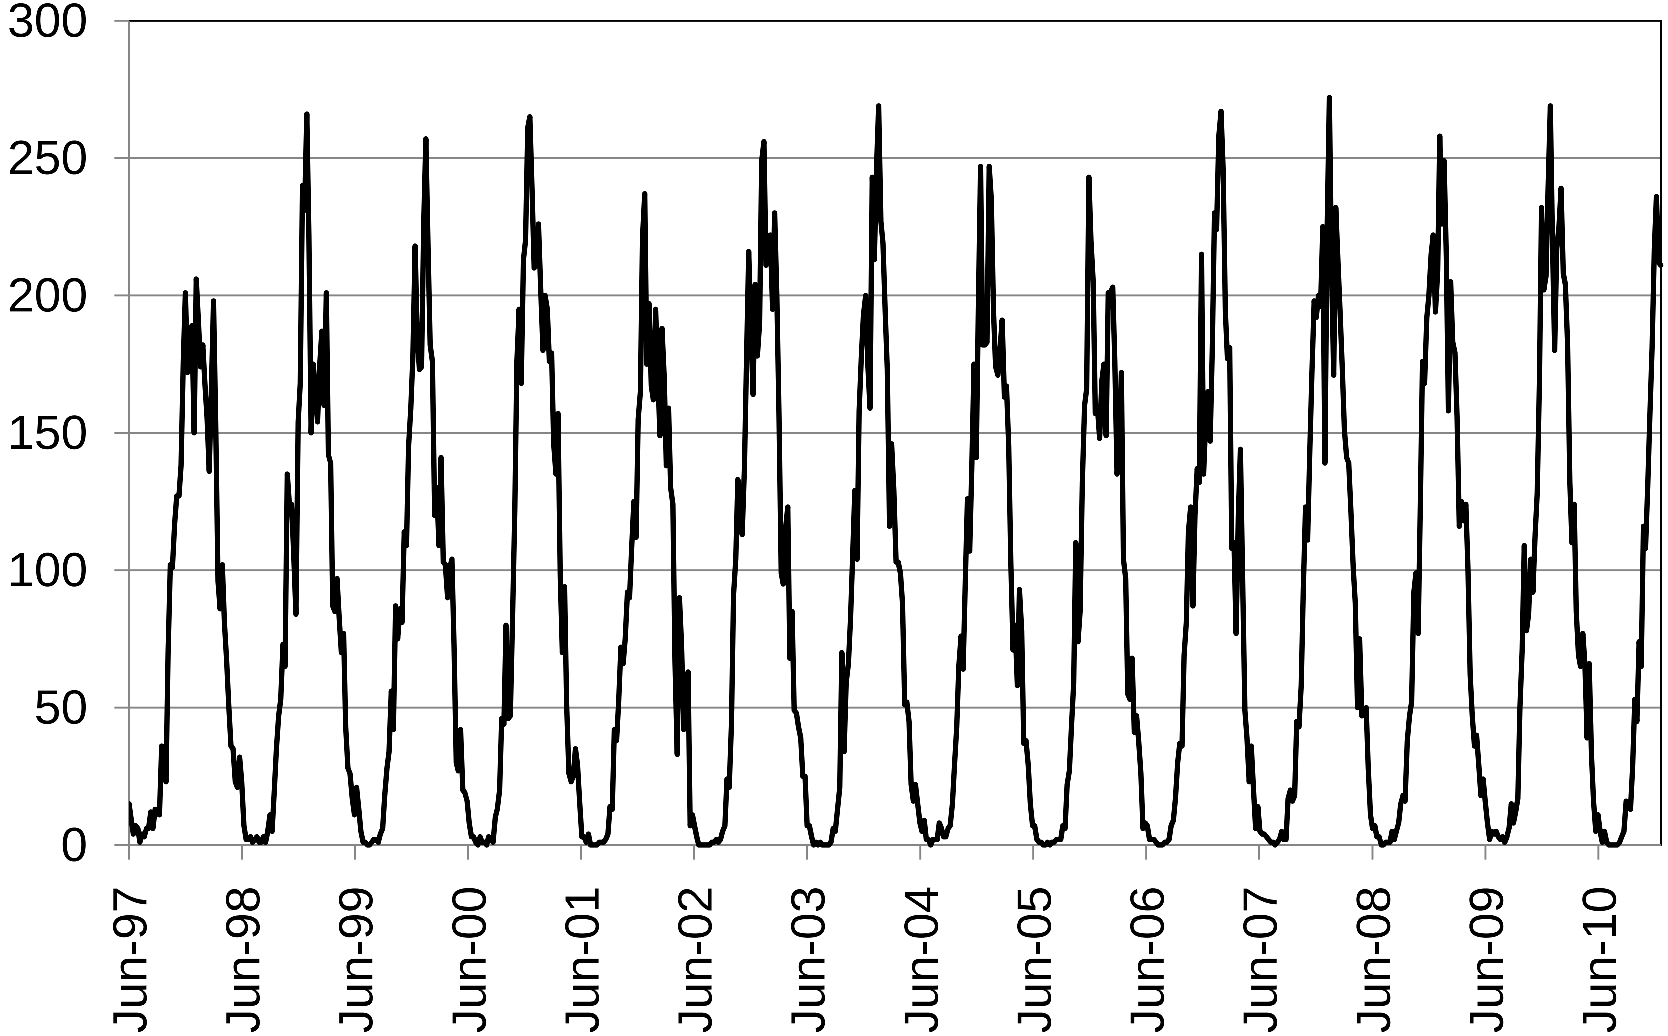

Figure A2. US Average Weekly Cooling Degree Days (CDD)

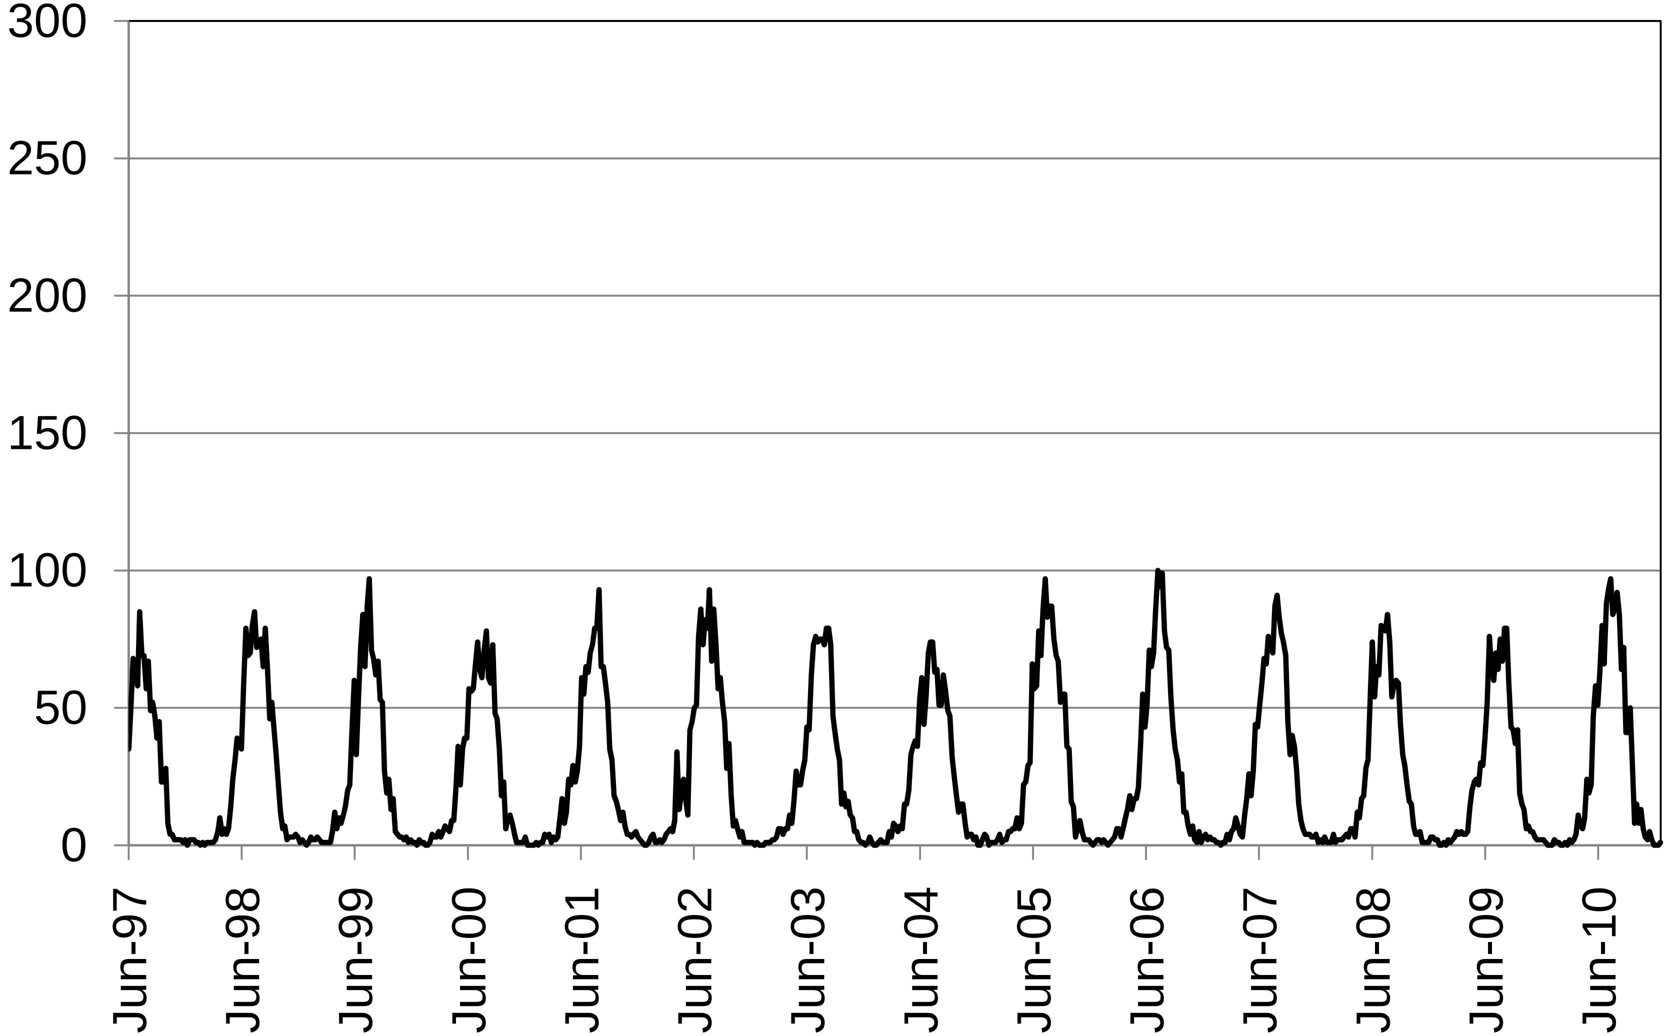

Figure A3. US Average Deviation from Normal Heating Degree Days (HDDDEV)

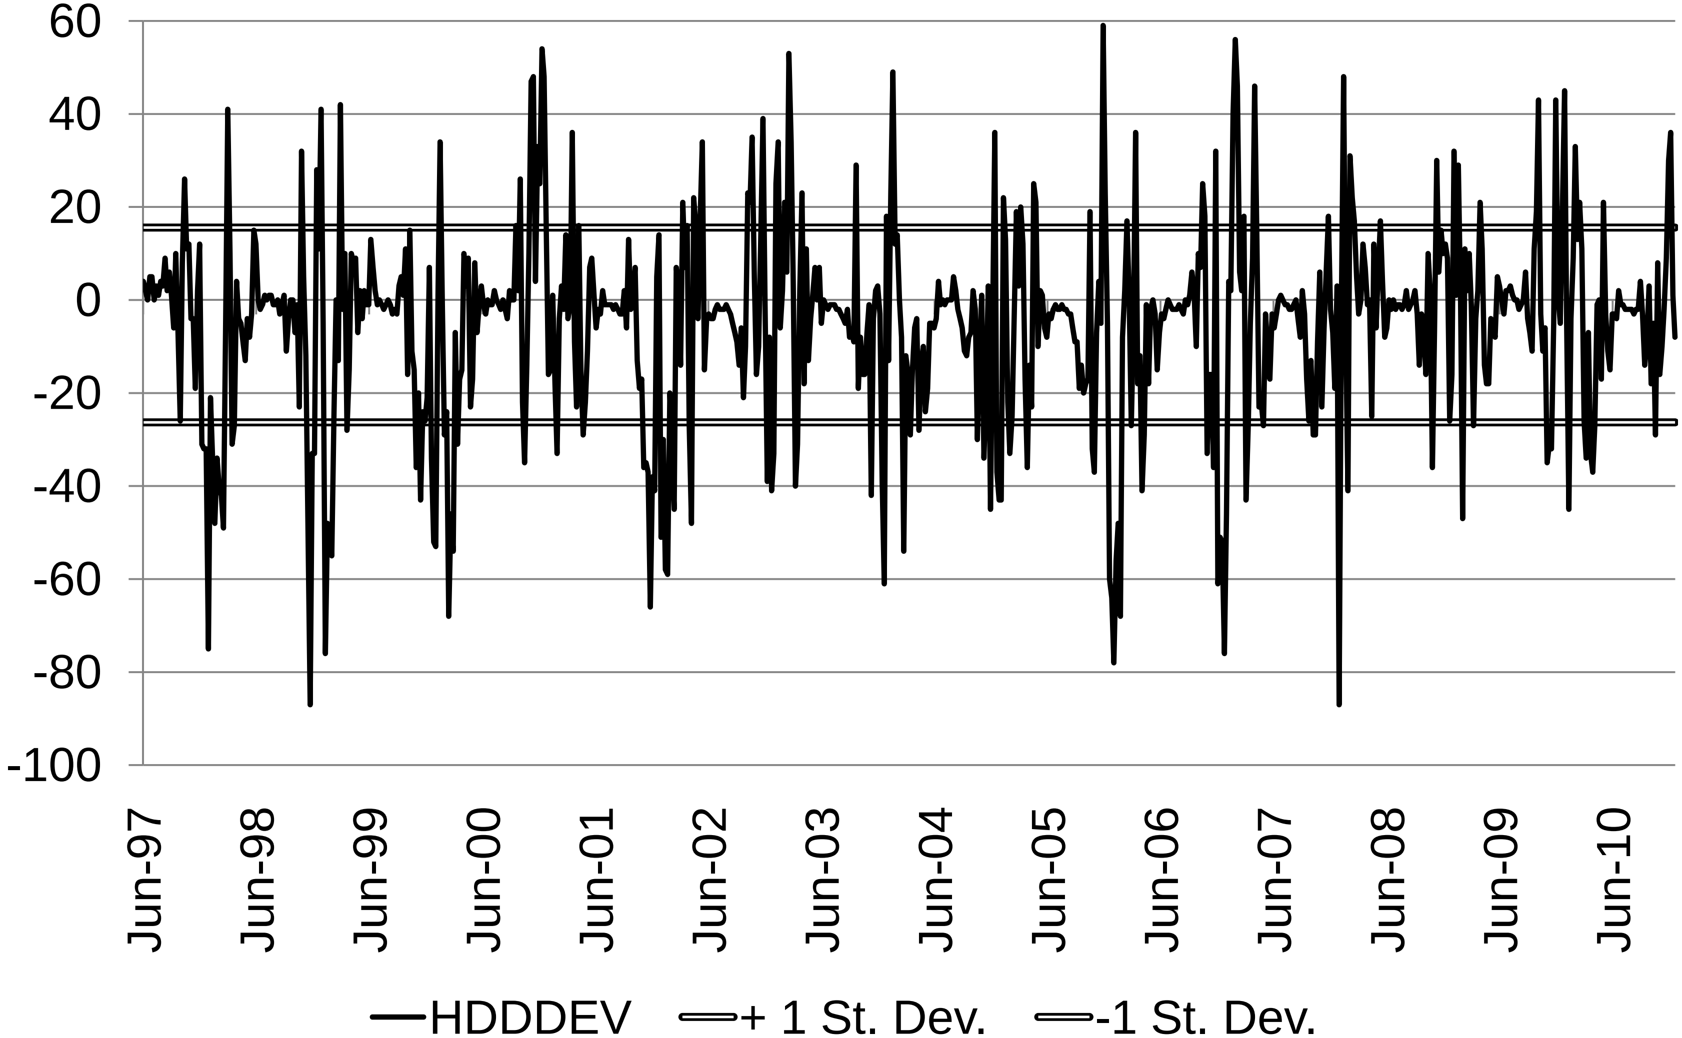

Figure A4. US Average Deviation from Normal Cooling Degree Days (CDDDEV)

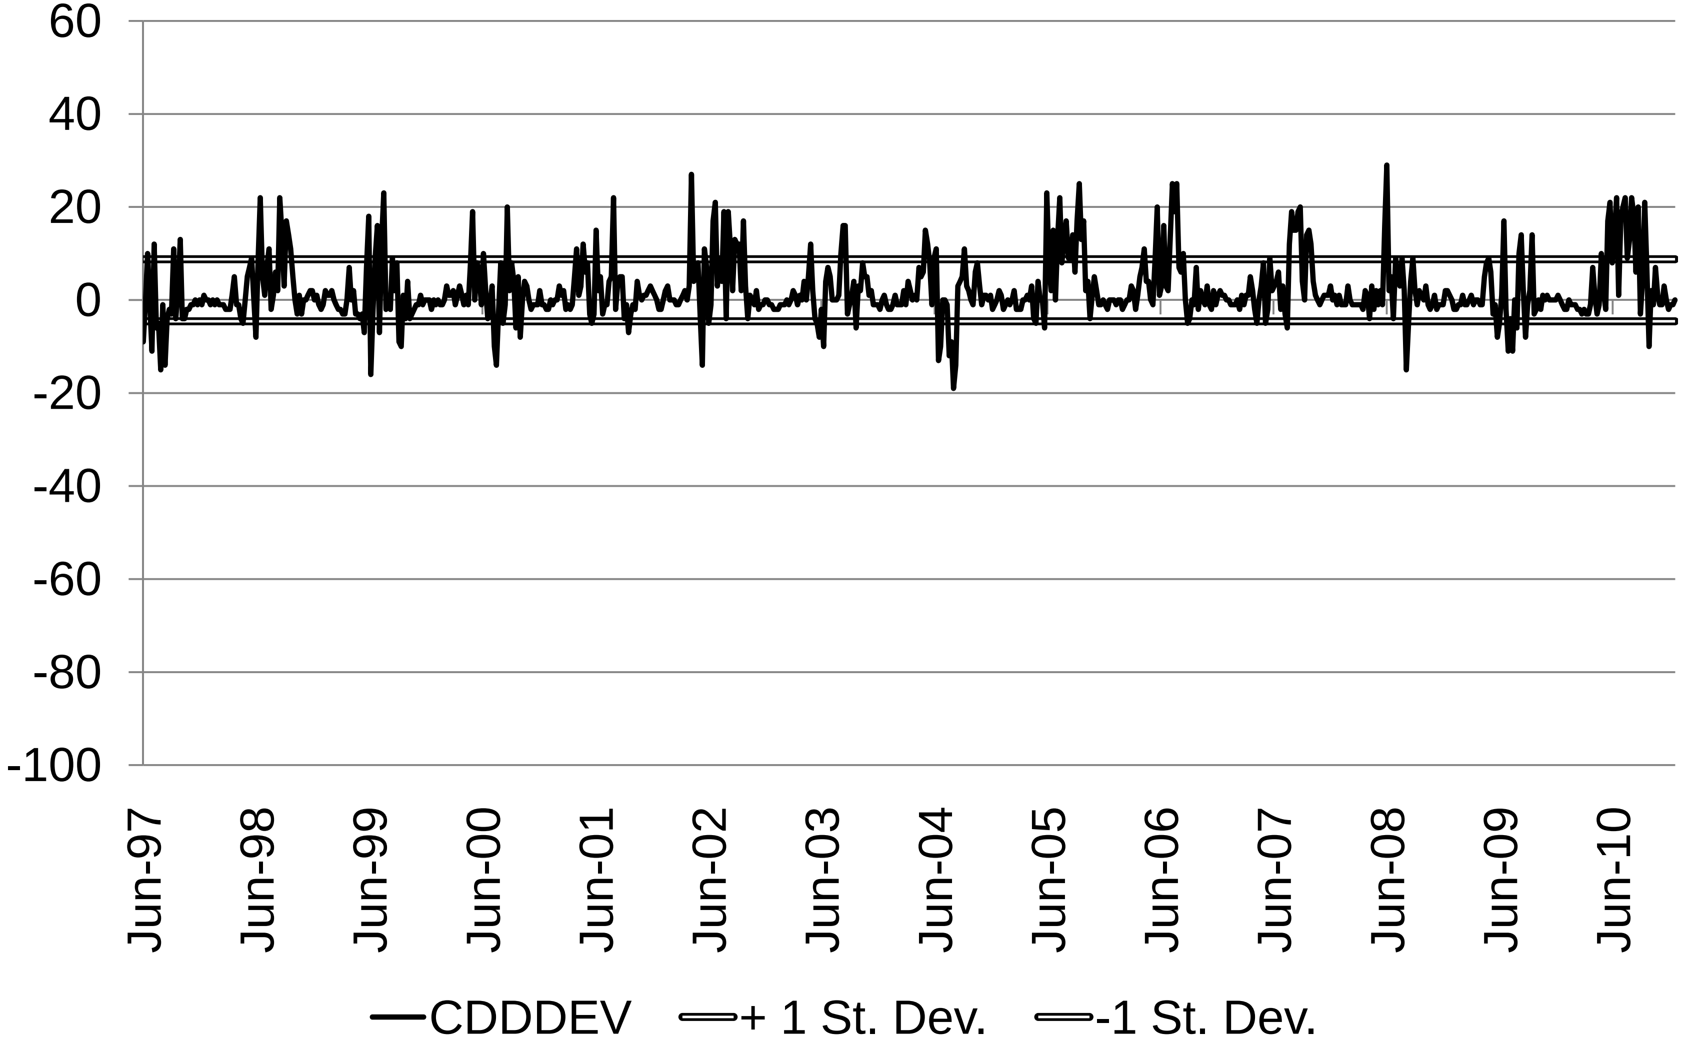

Figure A5. Shut-in Natural Gas Production Capacity in the Gulf of Mexico (SHUTIN), mmcf

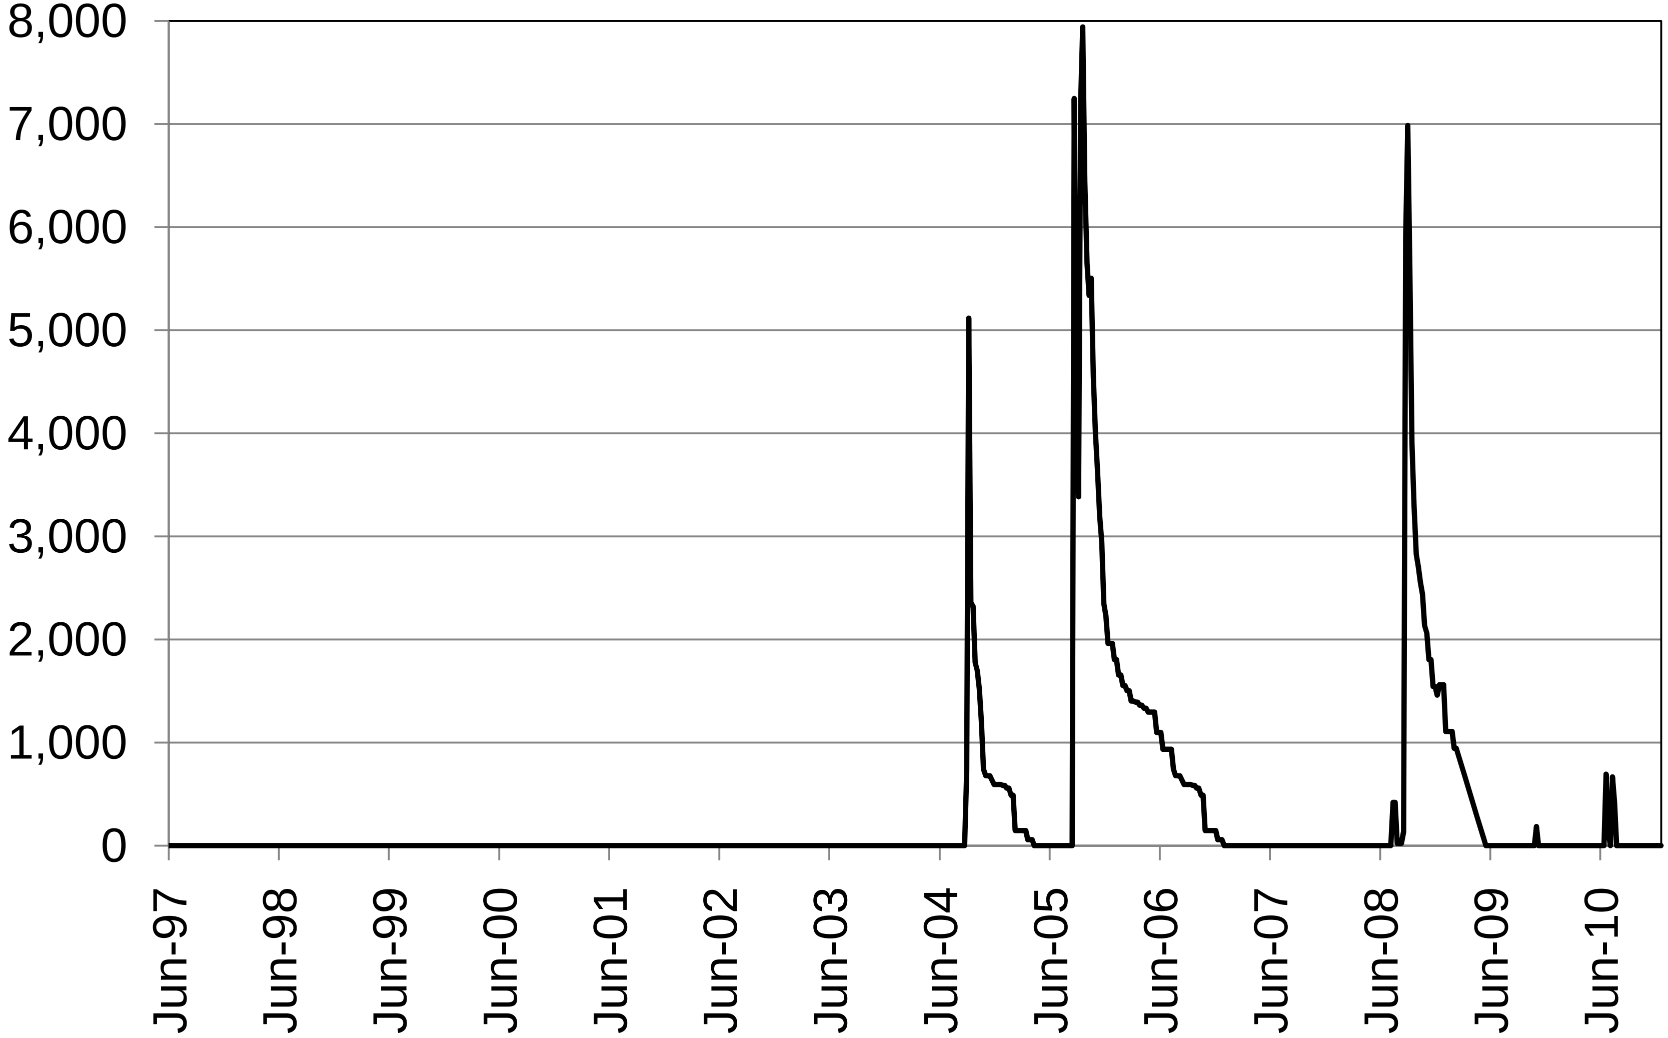

Figure A6. Differential from Running 5-year Average US Natural Gas Storage Levels (STORDIFF), Bcf