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*[**]

November 2010

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 prices 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 price of natural gas 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 are 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. for research on "Conversion of High-C/Low Value Feedstock into Fuels and Electricity: Exploring Technological Opportunities and How These Could Impact the Market."

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. Starting in December 2008, the price of crude oil started to recover from its low of $31.41/bbl. By February 2009 it had already risen to $44.76/bbl. During the same time, the price of natural gas continued to fall from its already low level of $5.37/mmBtu, dropping to $4.03/mmBtu in February. For most of 2009 the prices continued to diverge so that in October crude oil reached below $70/bbl while natural gas fell below $3/mmBtu. Natural gas then briefly spiked, but as of May 2010 crude oil was $86.19/bbl and natural gas was $5.37/mmBtu to $3.86/mmBtu. The ratio of the price of natural gas to the price of crude oil now stands near its lowest point in decades. So what happened to the strong tie between the prices that these studies have documented? Some believe that the recent price movements reflect a permanent rupture of the old tie between the two price series, arguing that fundamental changes in the

[1] The term cointegration is used to describe a certain relationship between two time series like the time series of oil and natural gas prices. This type of time series is 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 it is a single process (or combination of processes) that underlies the two time series, causing them to be non-stationary. When this is the case, the relationship between those time series 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).

marketplace have occurred. If so, then studies establishing cointegration are already outdated.

This is not the first time the natural gas price has appeared to decouple from the oil price. In the late 1990s and early 2000s, the price of natural gas was regularly *above* the level one might have predicted based on the historical relationship and the then prevailing price of crude oil. In an earlier era, the United States experienced a so-called gas bubble that had kept natural gas prices low relative to the then prevailing price of crude oil. Nevertheless, the statistical fact of cointegration appears to characterize the two time series throughout these periods of ups and downs. How is one to rationalize these seemingly contradictory facts? What do we really mean by cointegration if the average 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 analysts assert that the two prices have "decoupled." These assertions are often vague and open to alternative interpretations. "Decoupled" could mean any of the following:

- (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).[3] The results in the

[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.

[3] A Vector Error Correction Model (VECM) identifies and characterizes the cointegrating relationship(s) between two or more sets of time series data. The relationship is measured in the level of each time series variable. For both crude oil and natural gas, it also models the *change* in price for each as a function of *past* changes in the prices of both commodities. This is done to capture the effects of long-lived shocks to the price series. Price changes of both commodities are used in the model because we theorize that crude oil and natural gas prices are related to one another, so we would like to see whether changes in one commodity price provoke changes in the other. This second portion of the VECM is similar to a vector autoregression (VAR) model, except that a VAR only models the effect of the history of a single time series on itself. The other component is the error-correction mechanism (ECM, not to be confused with the VECM). The ECM measures the rate at which the time series returns to the cointegrating relationship after price shocks have caused it to deviate. So the VECM has three components: the cointegrating relationship (the long-run price relationship between the commodities), the effects of the lagged price changes (the shocks that push the commodity price away from its long-run relationship), and the ECM (the rate at which the commodity price is pulled back into its long-run relationship). The VECM models the relationship between crude oil and natural gas as if each commodity affected the other. But we theorize that crude oil prices should affect the natural gas price, not vice versa. The VECM identified the cointegrating equation and provided the ECM to measure the rate at which the natural gas price returns to its cointegrating relationship with the crude oil price. In order not to treat crude oil prices as being partially determined by natural gas prices, we turn to the conditional Error Correction Model (conditional ECM), which models the system as a VAR. Here, past price changes of Henry Hub natural gas prices affect the change in the natural gas price (pushing the gas price away from the cointegrating relationship). We add the ECM to measure the

four papers are broadly consistent with one another, although the details of the modeling and the parameter estimates vary. We report results based on our own modeling and tests along the lines of Brown and Yücel (2008), but include as well 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 formal specifications and the full set of statistical tests performed. However, these are detailed in the Appendix. Our conclusions are as follows.

Consistent with earlier studies, we find clear statistical evidence that the two series are cointegrated. Changes in the price of oil do appear to translate into changes in the price of natural gas. Movements in the price of natural gas away from the underlying co-integrating relationship tend to reverse, so that this relationship is re-established. However, recognition of the statistical fact of cointegration needs to be tempered with two additional points that we think have been insufficiently emphasized in the past literature.

First, there is an enormous amount of unexplained volatility in natural gas prices at short horizons. 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 will leave a large portion of the price of natural gas unexplained. The more statistically sophisticated approach of constructing a conditional Error Correction Model, which includes the cointegrating relationship, a set of exogenous explanatory variables, and accounts for the reversion of natural gas prices back to the

---

rate at which natural gas prices return to the cointegrating relationship found through the VECM, and also model the reaction that natural gas prices have to the concurrent change in the crude oil price. The conditional ECM thus conforms to our theory.

cointegrating relationship, still leaves a large portion of the volatility in natural gas prices unaccounted for.

Second, the cointegrating relationship, while clearly present in any period of the data, 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. In the 1989 to 2005 period, the price of natural gas seemed to be shifting up compared to the price of oil, but in recent years this reversed.

Therefore, although the two price series are 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 DATA

Figure 1 shows the raw spot price series for WTI crude oil and Henry Hub natural gas from 1991-2010 plotted together on the same graph.[4,5] 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. The eye can immediately detect some rough relationship between the

[4] 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.

[5] Both series are weekly day-ahead prices of commodities on the trade date 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.

two price series. The price spike of 2008 is the most dramatically clear example of this, but the price run-up from 2003 through 2007 also clearly reflects some tie. Even in the time period before 2002 this rough relationship seems to show up, though with less clarity.

Figure 2 shows the same two price series, but now reported in logs. In a graph of the absolute price levels, like Figure 1, where prices started off low and rose through time, a given percentage change in an early period may be unnoticeable, while the same percentage change in a later period appears very dramatic. The large scale demanded by later price movements causes earlier price changes to appear negligible. In logs, a given percentage price change appears the same size regardless of the price level. This is why using logged prices is a useful way to display and work with the data. Looking at Figure 1 it would be easy to conclude that natural gas prices have become more volatile, since there are several large spikes late in the series. But Figure 2 makes clear that, in percentage terms, there were comparably sized spikes earlier in the series, and there is no obvious suggestion that the series had become more volatile. Even the scale of the dramatic oil price spike of 2008 seems less unusual when reported in logs.

It is clear in Figure 2 that the natural gas price series is much more volatile than the oil price series. The annualized volatility of the log natural gas price series is 69%, 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.[6] Much of the volatility in the natural gas

[6] 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 will typically 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 will be the same.

price series appears to take the form of temporary spikes in the price. These spikes appear to have a relatively short duration. This will be key to the later discussion.

3. MODELING THE NATURAL GAS-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.[7] 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 and processing, and they serve different portfolios of end uses with only a modest overlap. 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.

It is worth emphasizing, in particular, that the crude oil and natural gas price series on which we focus are determined in different regional marketplaces. While the WTI crude oil price is for delivery in Cushing, Oklahoma and refers to a specific type of

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

oil produced in that region, it remains a benchmark price for crudes traded globally. The WTI price is tied to fluctuations in global demand and supply, moving relatively tightly with the prices of crudes delivered into other regions. 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. While the natural gas price in North America is linked to the fluctuations in supply and demand elsewhere in the globe, prices in different regions can move markedly apart from one another at times.

In fact, nothing like an energy content equivalence has been persistently observed. 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 other times as high as 19-to-1. Natural gas prices sometimes spike dramatically, without there being any noticeable change in crude oil prices. So is the price of natural gas tied to the price of oil, or not?

A casual reading of the industry press reveals no clear consensus on this issue. Some seem to take for granted that the two prices must maintain some rough parity with one another, although there seems to be uncertainty about where that parity level is properly drawn. The literature contains a variety of rules-of-thumb, including the simple 10-to-1 ratio, as well as more sophisticated burner tip parity rules. What is the formula that best describes the relationship?

Part of the problem is that a number of other variables have some short-run influence on the prices of both natural gas and oil. 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 in the United States is not seasonal, so the ratio of the prices must vary through the calendar year. Other variables that complicate the relationship are stochastic, but observable and readily incorporated as conditioning variables in a statistical analysis. 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. Finally, both the natural gas price and the oil price are subject to idiosyncratic shocks to supply and demand in their respective markets, shocks that briefly disturb the usual price relationship between the two and that are not easily translated into conditioning variables. To identify the underlying tie between the two commodity prices requires filtering out the effect of these various factors on the observed price relationship. This is a challenging task.

We break the 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 relationship. 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 are a few identifiable and recurrent outside factors – such as seasonality, episodic heat waves and cold waves and intermittent supply disruptions from hurricanes – that temporarily shift this fundamental tie in predictable ways. Third, there are a myriad of uncatalogued temporary disruptions to the supply and demand for natural gas that pull the price of natural gas away from the fundamental tie.

Fourth, when the natural gas price has been pulled away from the fundamental tie, it will predictably drift back towards it. The statistical analysis attempts to identify and filter out the three identifiable components, the first, second and fourth. Once each of these identifiable components of the movements in the natural gas price have been filtered out, we are left with the residual or unexplained shocks to the natural gas price.

The purpose of our analysis is to uncover these three identifiable components and describe them in an accessible way, and to determine how much of the volatility in the natural gas price is accounted for by them and how much is left unexplained. In the statistical analysis presented in the appendix the three components are modeled together in one system of equations, and the parameters are estimated simultaneously. However, in the discussion here in the body of the text, we focus one at a time on the separate components, beginning with the first: the underlying fundamental tie between the two prices.

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

In the beginning of Section 3, we reported one possible structure for this fundamental tie and that was an equality of price in terms of the energy content of the two commodities. This would predict a price of oil that is approximately six times the price of natural gas. Brown and Yücel (2008) and Hartley, Medlock and Rosthal (2008) each report a competing rule-of-thumb of 10-to-1, widely cited among players in the oil patch. Brown and Yücel (2008) also report two other rules-of-thumb known as burner tip parity rules. One is based on competition 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.[8] For example, at a \$10/bbl oil price, the residual burner-tip parity rule predicts a natural gas price of \$1.10/mmBtu, while the distillate burner-tip parity rule predicts a natural gas price of \$1.26/mmBtu. When the oil price is \$60/bbl, the residual burner-tip parity rule predicts a natural gas price of \$7.86/mmBtu, while the distillate burner-tip parity rule predicts a natural gas price of \$11.56/mmBtu.

Each of these four pricing rules is charted in Figure 3. Although Figures 1 and 2 show data through mid-2010, the empirical results reported in the remainder of this paper rely only upon data through February 2009. Therefore Figure 3 and those which follow show data through this shorter horizon. 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 3 also shows the actual combinations of crude oil and natural gas prices observed in our data series as a scatterplot. Each data point represents a different week's pair of prices, with the week's crude oil price determining the point's location along the x-axis, and the week's natural gas price determining the point's location along the y-axis. Note that the scale required to cover the full range of the oil price series means that there is a dense cluster of points at the lower end of the range. Therefore, in Figure 4 we reproduce the rule-of-thumb graphs

[8] 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 $P_{HH,t} = -0.80 + (120\% / 5.825) P_{WTI,t}$.

and the data, but focus only on the lower portion of the range of oil prices, i.e., those below $30/bbl.

Neither Figure 3 nor Figure 4 shows the time dimension to the observed prices. Figure 5 provides a time-series representation of the performance 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. One can see the periods in which one or another rule appears to have more or less closely approximated the actual relationship.

One can summarize the data in Figures 3, 4 and 5 as follows. Of the four rules of thumb, the residual burner-tip parity rule prices were on average closest to the actual natural gas prices over the full period, 1991-2009, with an average error in logs of -0.090, i.e. just slightly above actual prices. The 10-to-1 rule prices had the second smallest difference from the actual natural gas prices, with an average error in logs of 0.135, i.e. just below actual prices. The energy content equivalence rule prices were on average further away from the actual natural gas prices, with an average error in logs of -0.405, which was considerably higher than actual prices. The distillate burner-tip parity rule exhibited the largest difference from actual natural gas prices, with an average error in logs of -0.419, also considerably higher than actual prices.

We can also evaluate the performance of each rule of thumb from two different perspectives: year-by-year and within crude oil price ranges in $/bbl. Which rule most closely approximated actual natural gas prices varied considerably over the spectrum of observed oil prices and over the 1991-2009 time period. Of the four rules of thumb, the residual burner-tip parity rule prices were closest to the actual natural gas prices, on average, in 9 of the 19 years examined (1993, 1996-1997, 1999-2002, 2004 and 2009), and when WTI prices ranged from \$25-30/bbl and from \$40-65/bbl. On average, 10-to-1 rule prices were closest to the actual natural gas prices in 7 of the 19 years (1991-1992, 1994-1995, and 2006-2008), and when WTI prices ranged from \$15-25/bbl and over \$65/bbl. The energy content equivalence rule prices were closest to the actual natural gas prices in 2 of the 19 years between 1991 and 2009 (2003 and 2005), and when WTI prices ranged from \$30-40/bbl. The distillate burner-tip parity rule prices were closest to the actual natural gas prices only in 1998, when WTI prices were at their lowest level, below 15/bbl.

The graphs in Figure 5 clearly demonstrate that each of the rules suffers from the identical failing: there is much more volatility in the natural gas price over the 1991-2009 period than can be accounted for by movements in the oil price, no matter what relationship is posited. Although a given rule may be systematically biased, if we center the mean absolute errors we obtain a measure of the minimum variation that could conceivably be achieved by each. The centered mean absolute errors for the four rules of thumb are very close to one another, ranging from 0.310 for the 10-to-1 and Energy Content equivalence rules to 0.324 for the distillate burner-tip parity rule. In short, the unaccounted-for volatility is relatively constant regardless of the rule of thumb. The differences between the centered errors are much smaller than the differences between the (un-centered) average errors in logs of the various rules.

Figures 3 and 4 also show a graph of the cointegrating relationship derived from the vector error correction model described in more detail in the Appendix. The cointegrating equation, in logs, is:

$$log\ P_{HH} = -1.0493 + (0.7342 \times log\ P_{WTI}).$$

The cointegrating relationship is linear in the logged prices. Converted back into dollars, the log-linear relationship is a slightly concave curve. To illustrate, the cointegrating relationship predicts a natural gas price of $1.90/mmBtu when the oil price is $10/bbl. At $60/bbl, the predicted natural gas price is $7.08/mmBtu. Were oil prices to reach $150/bbl, the cointegrating relationship predicts a corresponding natural gas price of $13.87/mmBtu. The VECM estimates this relationship together with an estimate of the other exogenous factors that determine natural gas prices as well as the dynamic effects of past changes in the prices of each series. Because certain of the data for these exogenous factors is only available starting in 1997, the cointegrating relationship identified in the equation above and in Figures 3 and 4 is derived from the shorter time window of June 1997 through February 2009.

Figure 6 depicts the performance of the cointegrating relationship at predicting the natural gas price from January 1991 through February 2009. The prediction error, measured by subtracting the natural log of the predicted gas price from the natural log of the actual natural gas price, is presented as a scatterplot comparable to those shown in Figure 5. We use the cointegrating relationship from the equation above in stand-alone form to calculate the predicted gas price used for Figure 6. None of the other variables from the larger VECM analysis have been included, nor have any of the dynamics from the VECM been included. The centered mean absolute error for the cointegrating relationship is 0.333, which is approximately the same as for the rules of thumb discussed above.[9] This repeats the earlier observation that the natural gas price series is just too

[9] Note that the VECM was estimated for the shorter 1997-2009 period. Focusing just on the 1997-2009 price data, the centered mean absolute error for the cointegrating relationship is 0.259. For the rules of thumb over this shorter window it ranged from 0.284 for the 10-to-1 and the energy content equivalence rules to 0.298 for the residual burner tip parity rule, to 0.321 for the distillate burner tip parity rule.

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 conditional ECM attempt to provide this accounting, and we now turn to examine how successfully they do so.

Conditional Variation

Typical Seasonality

Figure 7 illustrates the estimated impact of predictable seasonality on the natural gas price. This is based upon a typical seasonal weather pattern between 1997 and 2009, meaning a typical pattern of Heating Degree Days and Cooling Degree Days. Using the coefficients from our conditional error correction model, and holding constant the crude oil price and all other exogenous variables, we simulated the path of the natural gas price through the typical annual cycle. Figure 7 is centered on the average price level, which occurs in the third week of June and again in the first week of December. From this benchmark, the price peaks at about 113.5% in March. The trough is at 86.3% in September. The total amplitude of the seasonal variation in the natural gas price is 27 percentage points. At a base price of $7/mmBtu (the June and December price), this is a range of $1.90/mmBtu.

Figure 8 compares the seasonal variation around the cointegrating relationship to the scale of the observed prediction errors in natural logs of the prices. 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 1991-2009 is 0.333. The seasonality coefficient, however, only ranges as high or low as +/- 0.125, 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 19% of natural gas volatility.

Unseasonably Cold or Warm Weather

Unseasonably cold or warm weather is measured by deviations from the normal number of either Heating Degree Days (HDD) or Cooling Degree Days (CDD). To illustrate the impact of these weather variables, we use the estimated coefficients in the conditional ECM as reported in the Appendix, and construct the impulse response function shown in Figure 9. The impulse employed in the figure is a typical cold spell lasting two weeks, with the first week exhibiting an HDD level 20 degree days above the average and the second week exhibiting an HDD level 15 degree days above the average.[10] The figure shows the resulting path of the natural gas price as estimated by the conditional ECM. This causes a deviation in price that peaks at 3.4% above the base level. To scale this, suppose that the oil price were at $50/bbl, and assume normal seasonal conditions; then the predicted natural gas price in the second week of March would be $6.71/mmBtu. Given this abnormal seasonal cold spell, the price is predicted to peak at $0.23/mmBtu higher than the predicted price under normal conditions. Thereafter, the price gradually falls back. The half-life of the price impact from the cold spell is approximately 7 weeks.

An unseasonable warm spell in the summertime is reflected in the deviation from CDD variable. In our dataset, a typical warm spell lasts only one week and on average

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

involves a 10 degree day variation. The heat wave prompts an impulse response of 3.3% above the benchmark natural gas price. The price that would prevail under normal seasonal patterns in the first week of August when the oil price is held at $50/bbl is $5.33/mmBtu. The warm spell produces a peak impact on the natural gas price of $0.18/mmBtu in the first week of the temperature deviation.

Figures 10 and 11 show that neither type of shock, on average, generates a very large impact on natural gas prices. In order to illustrate the small magnitude of the cold snaps and heat waves, we have plotted the actual natural gas price series in Figures 10 and 11 as a dashed gray line. We plotted actual prices that occurred after episodes in which HDD and CDD deviations in the data set matched our impulse-response experiments. Note that Figures 10 and 11 plot actual prices against the unexpected weather impulse response. The effects of other exogenous variables are included in these actual prices, so not all of the price variation represents unexplained volatility. A statistical test of the contribution of the exogenous variables to actual natural gas price volatility is included in the Appendix. The purpose of plotting actual prices against the impulse response in Figures 10 and 11 was to illustrate the relatively minor impact of the abnormal weather variables to natural gas price changes in the conditional ECM. We do not attempt to remove the effects of all exogenous variables in these figures, which would represent all of the unexplained volatility in natural gas prices. The actual price movements dwarf the movements predicted by our model's coefficients, and the effect of the HDD and CDD deviations is not even noticeable after the impulse. Where the cold snap was predicted to increase prices by 3.4% the week after initial impact, the actual prices fell by 0.56%. Where the heat wave predicted an instantaneous 3.3% price increase, actual prices fell 30%. This illustrates the relative contribution of this exogenous conditioning variable to the volatility of the natural gas price.

Although the impulse responses of HDD and CDD shocks are similar in form, the specific values are different. First, the typical scenario is a shorter spike in CDD than in HDD. Second, the coefficient on the deviation from CDD is larger than the coefficient on the deviation from HDD. As shown in Figure 9, the price after a CDD shock immediately begins to fall back to its original level with a half-life (identical to that of the HDD shock) of about 7 weeks.

Shut-in Production Due to Hurricanes

Figure 12 plots the impulse response function on prices from a shut-in of natural gas production as a result of a hurricane. We construct it using the pattern of shut-in that resulted from Hurricane Ivan in September of 2004. At its peak, Hurricane Ivan shut-in over 5.15 million mmBtu of natural gas production, representing what was estimated at the time to be over half the gas production capacity in the Gulf.[11] The disruption lasted over 33 weeks, and the average shut-in during this time was 836,768 mmBtu. The average shut-in production level represented a curtailment of about 7.4% over the 33-week period, which lasted until mid-April the following year. To plot the impulse response, we used the coefficients from our conditional ECM, and held constant the crude oil price and all other exogenous variables. The precise time pattern of the disruption is graphed with the dashed line in Figure 12 measured against the right-hand-axis. The model's estimated impact on price for a supply disruption with this profile is graphed with the solid line in Figure 12 measured against the left-hand-axis. As seen in

[11] Revised data from the Energy Information Administration put the actual shut-in production capacity against August 2004 production levels at 47%.

the figure, the maximum impact of the shut-in is a price premium of about 5% over a seasonally adjusted benchmark price of $5.51/mmBtu. This is an impulse response of $0.27/mmBtu over the benchmark. The half-life of the shut-in itself is 3.4 weeks, and the half-life of the price impact is 34 weeks.

Figure 13 superimposes the actual Henry Hub natural gas prices over the same period. The model's estimated effect is much smaller than actual price movements over the period, which were also much more volatile than our set of exogenous variables were able to account for. This reveals the relative magnitude of the unaccounted-for volatility.

Storage Levels

Finally, we analyze the impulse response of deviations from normal storage levels on the natural gas price. Storage differentials arise and dissipate over a longer period of time than unseasonable weather. The dashed line in Figure 14 charts the storage differential that arose between January and October of 2003 and is measured on the right-hand-axis. We used this sample of actual data rather than constructing an experiment due to the fact that the storage differential in this period almost perfectly matched the median storage differential over the entire 1997-2009 period in both duration and degree. The differential was not likely due to any shock, but rather simply a deviation from the average storage level over our period of study. The average shortfall over the 38-week episode was about 360 million mmBtu, although at its maximum the shortfall was nearly 629 million mmBtu. Figure 14 also charts the model's predicted price impact with a solid line. The maximum price impact was reached on the 22nd week of the disruption with a price premium of 15% over the seasonally-adjusted base price of \$6.30/mmBtu.[12] This impulse response corresponds to a price increase of \$0.93/mmBtu over the benchmark. The model's estimated half-life of the impulse response to the storage shortfall from its peak was about 23 weeks. The impact of a supply overhang is roughly the mirror image of the supply shortfall.

Figure 15 superimposes the actual price movements of Henry Hub natural gas onto Figure 14. As in the case with shut-in production, the actual price movements dwarf the impulse response estimated by the model. In the real world, there was more than simply the storage differential variable in play. However, even taking this into account, the model is unable to explain most of the price volatility in natural gas that actually occurred between January and October of 2003.

Disturbances to the Relationship and the Rate of Recovery

Even after accounting for the variation in the price of natural gas that is ascribed to fluctuations in the price of oil as well as the variation in the price of natural gas that is ascribed to seasonality and the identifiable exogenous disturbances already isolated, there is a large amount of unexplained volatility in the natural gas price. This is presumably attributable to many uncatalogued events and evolving changes in the natural gas marketplace that are not reflected in these exogenous variables.

These unexplained movements are events that pull the price of natural gas temporarily away from the identified relationship with the price of crude oil. Gradually, the identified relationship is reestablished. How quickly one can expect this to occur is

[12] A storage shortfall is likely to arise coincident with other events, such as a shut-in of production or a severe weather event. However, in Figure 14 we examine the impact of a sudden storage shortfall arising in mid-winter, one that is resolved quickly, and which arises in isolation.

shown in Figures 16 and 17. There is a coefficient in each of the error correction models that measures the rate at which the natural gas price returns to the cointegrating relationship, both after changes in the oil price and after deviations in the natural gas price from its cointegrating relationship with the oil price. This is the error correction term, and the next paragraphs illustrate its mechanics.

Figure 16 depicts the estimated behavior of natural gas prices when the crude oil price rises 20%, from $50/bbl to $60/bbl, and all other variables are held constant. The dotted line represents the target price of natural gas based on the level of the crude oil price. The solid line shows the path of the natural gas price as it seeks the new equilibrium over a span of weeks. The half-life of this movement is nearly 11 weeks.

Figure 17 shows the estimated behavior of the natural gas price after it is subjected to a shock that is not due to any of the variables that are included in the model. That is, the crude oil price and all of the conditioning variables are held constant. In this case, the natural gas price spikes up by 162%, or $10/mmBtu from its baseline of $6.19/mmBtu when the crude oil price is steady at $50/bbl. Under this scenario, the half-life for the return of the natural gas price to its cointegrating relationship is nearly 8 weeks.

Explained and Unexplained Volatility

The implementation of the VECM and conditional ECM modeling techniques resulted in general improvements over the rules of thumb in fitting the natural gas price series given the price of oil and the values of our six exogenous variables. The centered mean absolute errors are much smaller than any of the rules of thumb. For the conditional

ECM, the centered mean absolute error was 0.202, while the smallest centered mean absolute error of the rules of thumb for the 1997-2009 period was 0.284, shared by both the energy content equivalence and the 10-to-1 rules.

Despite the considerable improvements in fit, a large amount of the volatility in natural gas prices could not be explained by the crude oil price and other exogenous variables. The portion of the volatility in the natural gas price explained by the combination of exogenous variables in the conditional ECM is approximately 16%. That means the fraction of variance in natural gas prices unexplained by our model is nearly 84%. This result led us to explore possible reasons for this weak explanatory power.

4. A CHANGING RELATIONSHIP OVER TIME?

Has the fundamental relationship between natural gas prices and oil prices been changing over time? Have the prices decoupled? If so, in what sense? Have they permanently broken away from the old relationship and moved into a new relationship, or have they temporarily broken away from the usual relationship, only to return at a later date? Or has the relationship between crude oil and natural gas prices become completely broken, with no sign of any pricing relationship whatsoever?

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.[13] They go a step further to specifically attribute this increase to the increased demand for natural gas arising from the installation of advanced CCGT power plants with significantly improved heat rates.

Figure 18 illustrates the dramatic shift in the Villar and Joutz predicted cointegrating relationship between crude oil and natural gas prices from the beginning of their data set in 1989 to the end in 2005. The trend curves are circumscribed by the actual price range of WTI crude contained in the 1989 to 2005 monthly price dataset. The authors apparently captured a time period in which natural gas prices were rising relative to crude oil prices. In our dataset, covering 1997-2009, we find that the cointegrating relationship shifted downward slightly, so that the level of natural gas prices fell relative to the price of crude oil. This is despite substantial overlap.[14] Given that the 1997-2009 cointegrating relationship we uncovered lies between the 1989 and 2005 Villar-Joutz cointegrating relationships, something must have changed between crude oil and natural gas markets such that the natural gas price became relatively less expensive. However, despite this shift in the level, the evidence of a relationship between the two time series remains strong.

To further understand the stability of the relationship, we split our dataset into two segments and re-estimated our system of equations separately on each segment. The first segment spans the weekly data from June 13, 1997 to December 16, 2005. The second segment covers December 23, 2005 through February 20, 2009. The results are that the

[13] As mentioned earlier, according to Hartley, Medlock and Rosthal (2008), the tie between natural gas and crude oil prices is not direct, but mediated by residual fuel oil prices.

[14] Villar and Joutz modeled a rising trend in natural gas prices over their sample period using a constant coefficient on time. This captures the in-sample trend. They cautioned against erroneously extrapolating this trend beyond their sample period.

relationship has shifted, modestly, but a fundamental relationship continues to exist. The two different cointegrating relationships are:

$log\ P_{HH}$ = -2.0142 + (0.986 x $log\ P_{WTI}$) for the 1997-2005 period, and

$log\ P_{HH}$ = -0.8598 + (0.6782 x $log\ P_{WTI}$) for the 2006-2009 period.

Figure 19 shows these two lines against a backdrop of the raw natural gas price and oil price combinations that separately lead to the two estimated equations. The line for the 1997-2005 segment is somewhat less steep than the Villar and Joutz December 2005 cointegrating relationship, as one might expect given that the Villar-Joutz cointegrating relationship is a snapshot at the end of our 8-1/2 year segment. The 2006-2009 segment lies somewhat closer to the Villar and Joutz January 1989 cointegrating relationship.

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 1997-2009 period as a whole. The 1997-2005 model accounts for nearly 21% of natural gas volatility through the crude oil price and the included conditioning variables. The 2006-2009 model accounts for nearly 31% of the price volatility in natural gas. These are considerable improvements to the 16% achieved by the original conditional error correction model. Nevertheless, there remains a large amount of unexplained volatility in the natural gas price even in each of these separately estimated time windows.

In contrast to the Villar-Joutz cointegrating relationships, the segmented cointegrating relationships we conducted do not assume a continuously increasing trend in natural gas prices relative to oil prices. Our modeling exercise suggests that natural gas actually shifted from one historical relationship with crude oil to one in which natural gas was relatively discounted to crude oil. This shift appears to have occurred some time between 2005 and 2009.

While there are clearly two distinct lines representing cointegrating relationships, keep in mind that the division of the data was arbitrary. Figure 19 could simply be depicting two snapshots of a gradually shifting relationship over time, and our separation of the data into these segments may have simply reflected the partial extent of this shifting relationship. Therefore we do not try to identify a structural break like those identified by Dvir and Rogoff (2009) in the much longer crude oil price time series.

The contrast between our findings and the Villar and Joutz results (which cover data only through the end of 2005) reinforce the hypothesis that the crude oil-natural gas price relationship is capable of shifting, and that it likely did so over the course of the 1997-2009 period we examined.

CONCLUSION

Consistent with industry assumptions, we have shown that there is a statistically significant pricing relationship between crude oil and natural gas. Use of both vector error correction and conditional error correction model formats has enabled us to characterize this relationship.

We have also been able to characterize the path of natural gas prices as they return to the pricing relationship with crude oil after gas price shocks. Including conditioning variables related to weather and supply issues concerning natural gas also enabled us to identify the mechanics of the seasonal fluctuation in natural gas prices, and to characterize the behavior of natural gas prices in response to hurricanes, heat waves and cold waves, and deviations from normal supply levels. However, despite these insights, our model of the 1997-2009 period only accounts for about 16% of the volatility in natural gas prices. There are many events that cause a price response in natural gas that are not accounted for in our modeling exercise.

The question of whether the price series of crude oil and natural gas have decoupled, and in what sense, required that we examine whether the price relationship had somehow changed over the course of our 1997-2009 data window. By breaking the data into distinct segments, we could determine whether more than one relationship existed, or indeed if a more recent segment failed to establish any cointegrating relationship whatsoever. This would allow us to ascertain whether any of the three interpretations of decoupling that we posited in the Introduction had occurred.

To test the possibility of a changing price relationship between crude oil and natural gas, we segmented the data into a 1997-2005 period and a 2006-2009 period and ran the same modeling exercises. We also compared these results with selected results from a Villar and Joutz paper that covered the 1989 through 2005 period. The results showed that between the two periods there are distinct relationships between the commodities, though whether there are only two relationships, or whether the difference represents a gradual shifting relationship over time, are not ascertained. The performance of the modeling exercise across the two periods does improve the ability of the model to account for the volatility in natural gas prices, however. The conditional error correction model corresponding to the 1997-2005 period can account for about 21% of natural gas volatility, while the model for the 2006-2009 period accounts for about 31%.

To address the question of decoupling in the first sense: yes, natural gas prices often break away from their usual relationship with crude oil prices, and to a large degree. However, in the 1997-2009 period we studied, these “decoupling” episodes were temporary, and the natural gas price always returned to some stable relationship with the crude oil price. In the second sense of “decoupling,” our results suggest that the relationship between oil and natural gas prices has probably experienced at least a gradual shift over the 1997-2009 period. This is evidence that, at least once over the 1997-2009 period, an initial crude oil-natural gas price relationship was permanently broken and a new equilibrium relationship was formed. As for the last interpretation of decoupling – that the two price series no longer maintain any relationship with each other at all – we find no evidence that this has occurred. Despite the appearances to the contrary, the statistical evidence of a cointegrating relationship between the two fuels, at least when measured from the end of December 2005 through February of 2009, has strengthened.

In short, despite large temporary deviations, natural gas prices continue to exhibit evidence of a cointegrating relationship with crude oil prices, and gas prices consistently return to a long-run relationship. However, this relationship has apparently shifted at least once over a 12-year period to a new equilibrium. There is no statistical evidence to support the claim that a relationship between the two price series has been completely severed.

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.

Dvir, E., and Rogoff, K., 2009, The three epochs of oil, National Bureau of Economic Research Working Paper Series, Working Paper 14927, 59 pp.

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.

Figure 1. The Natural Gas and Crude Oil Spot Prices, 1991-2010.

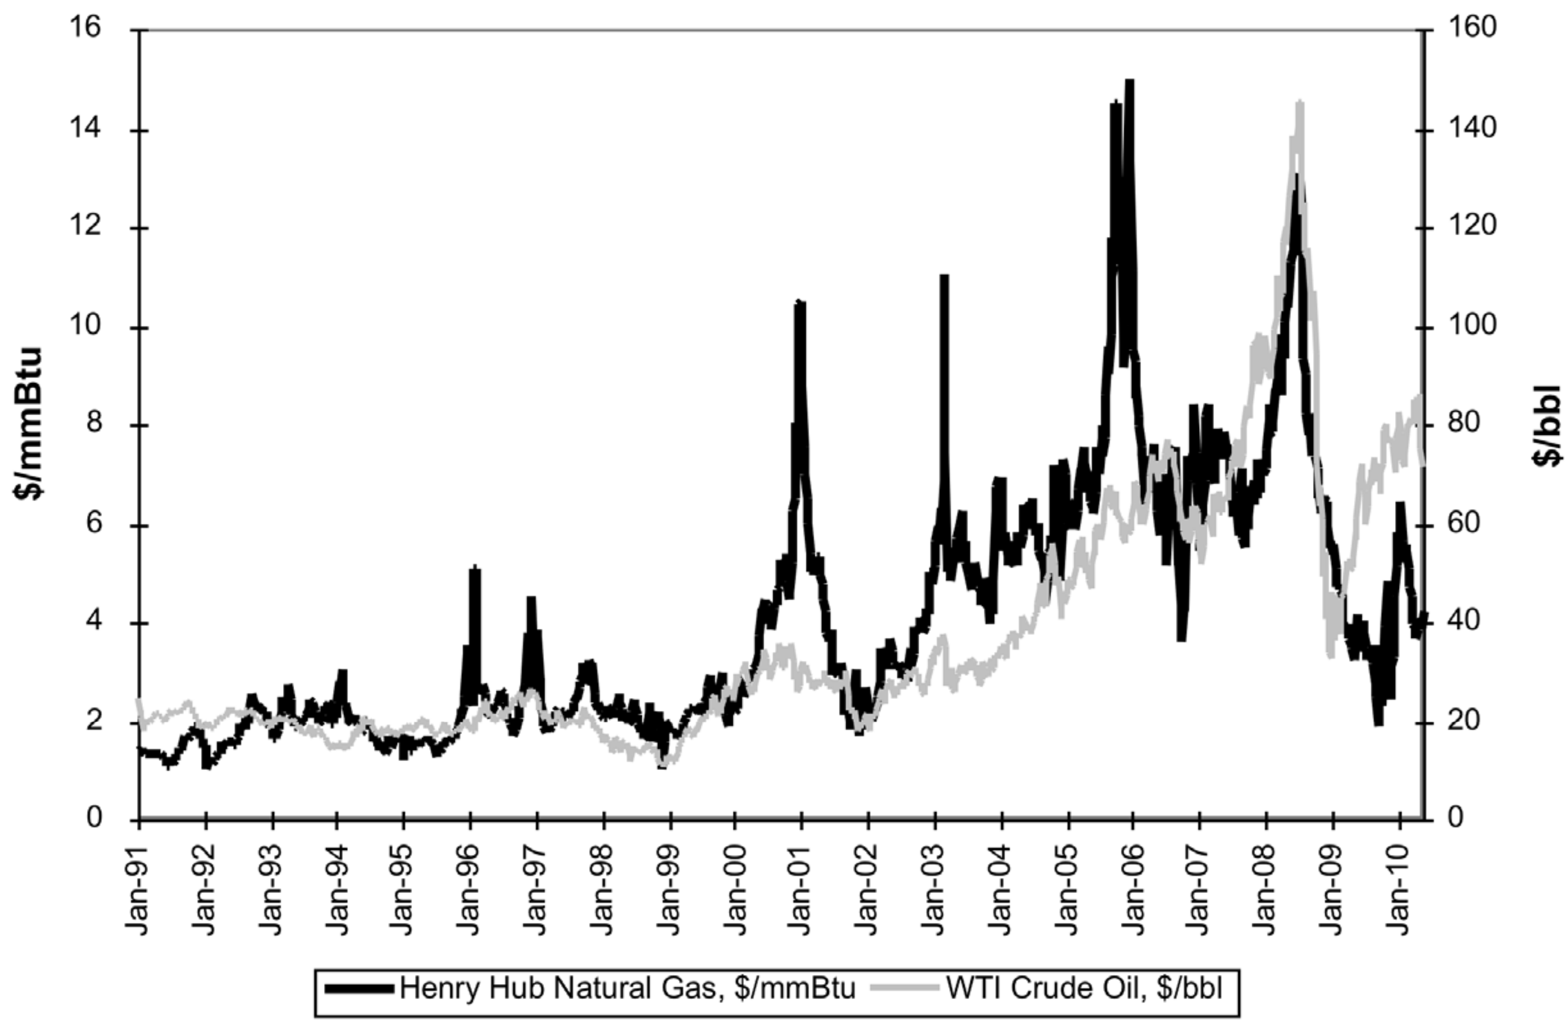

Figure 2. Log Values of the Natural Gas and Crude Oil Spot Prices, 1991-2010.

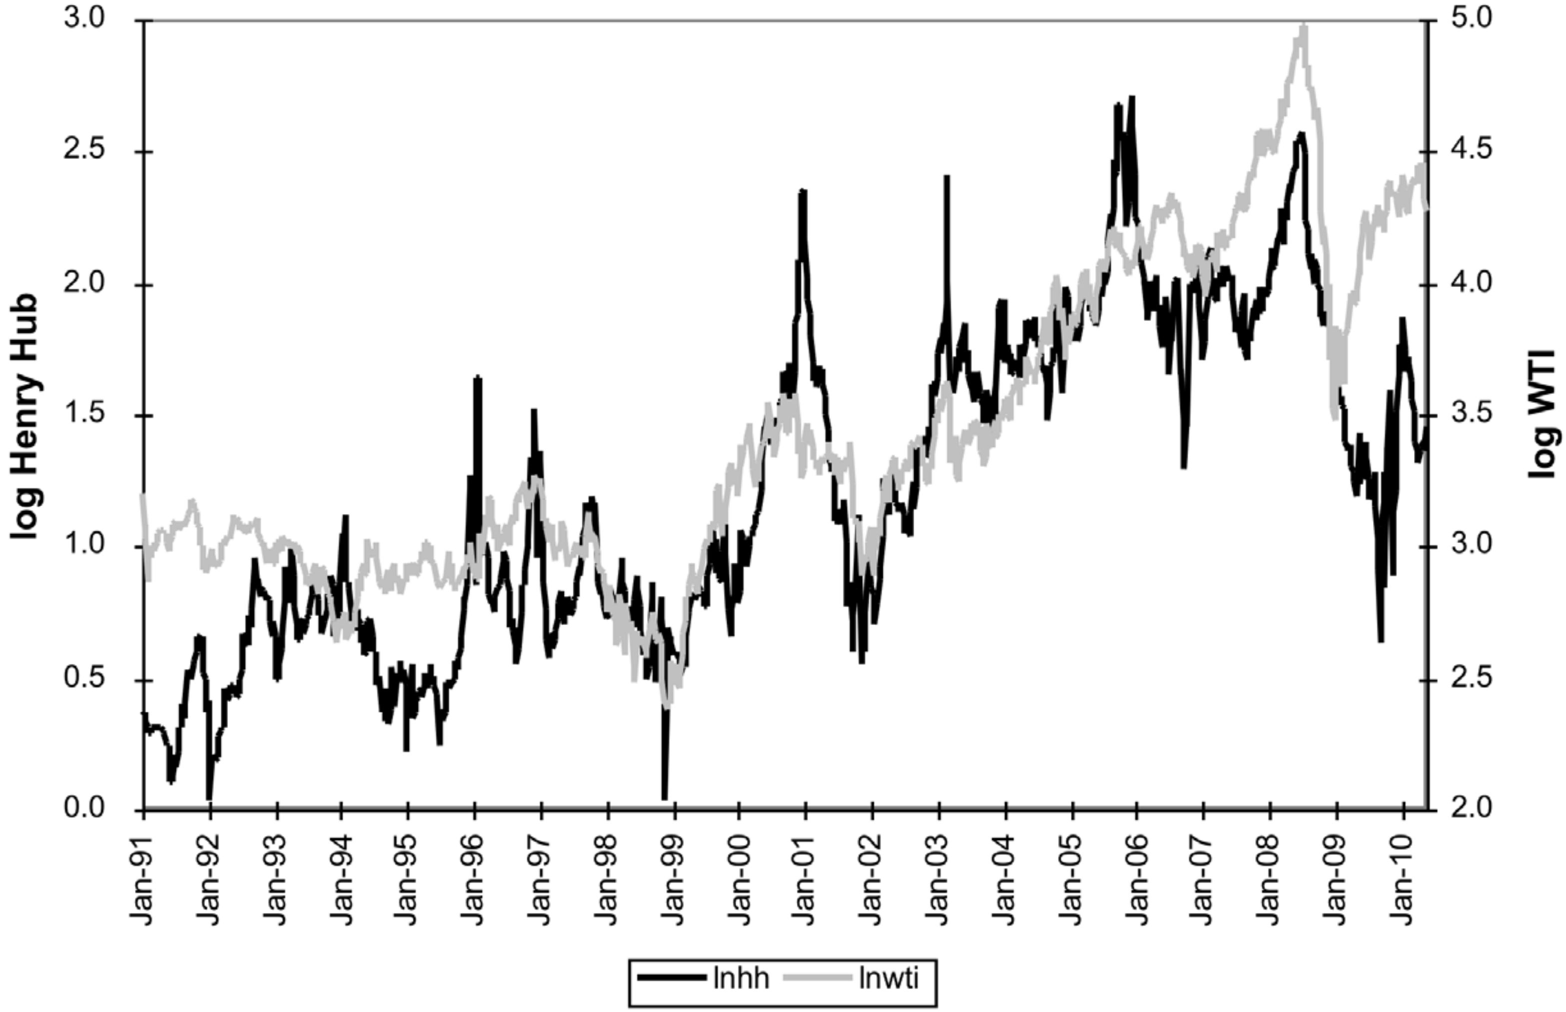

Figure 3. Price Benchmarks Versus Observed Prices
(Full Range, 1997-2009)

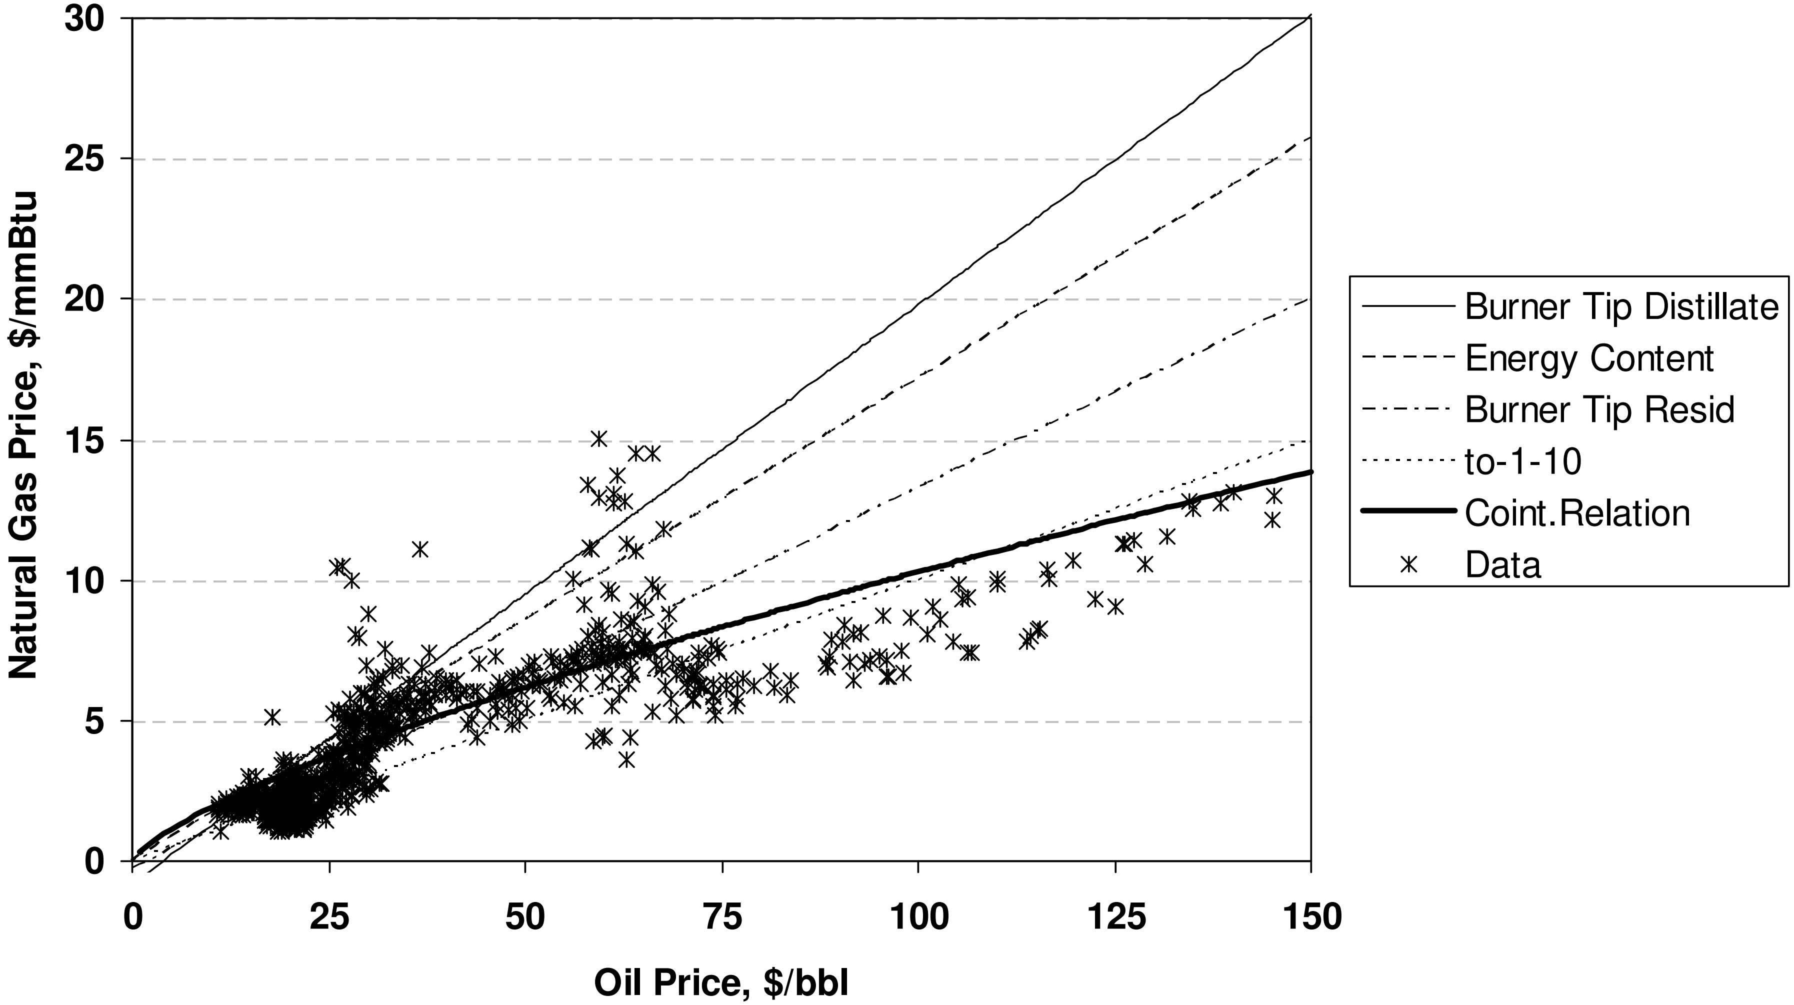

Figure 4. Price Benchmarks Versus Observed Prices (Low Range)

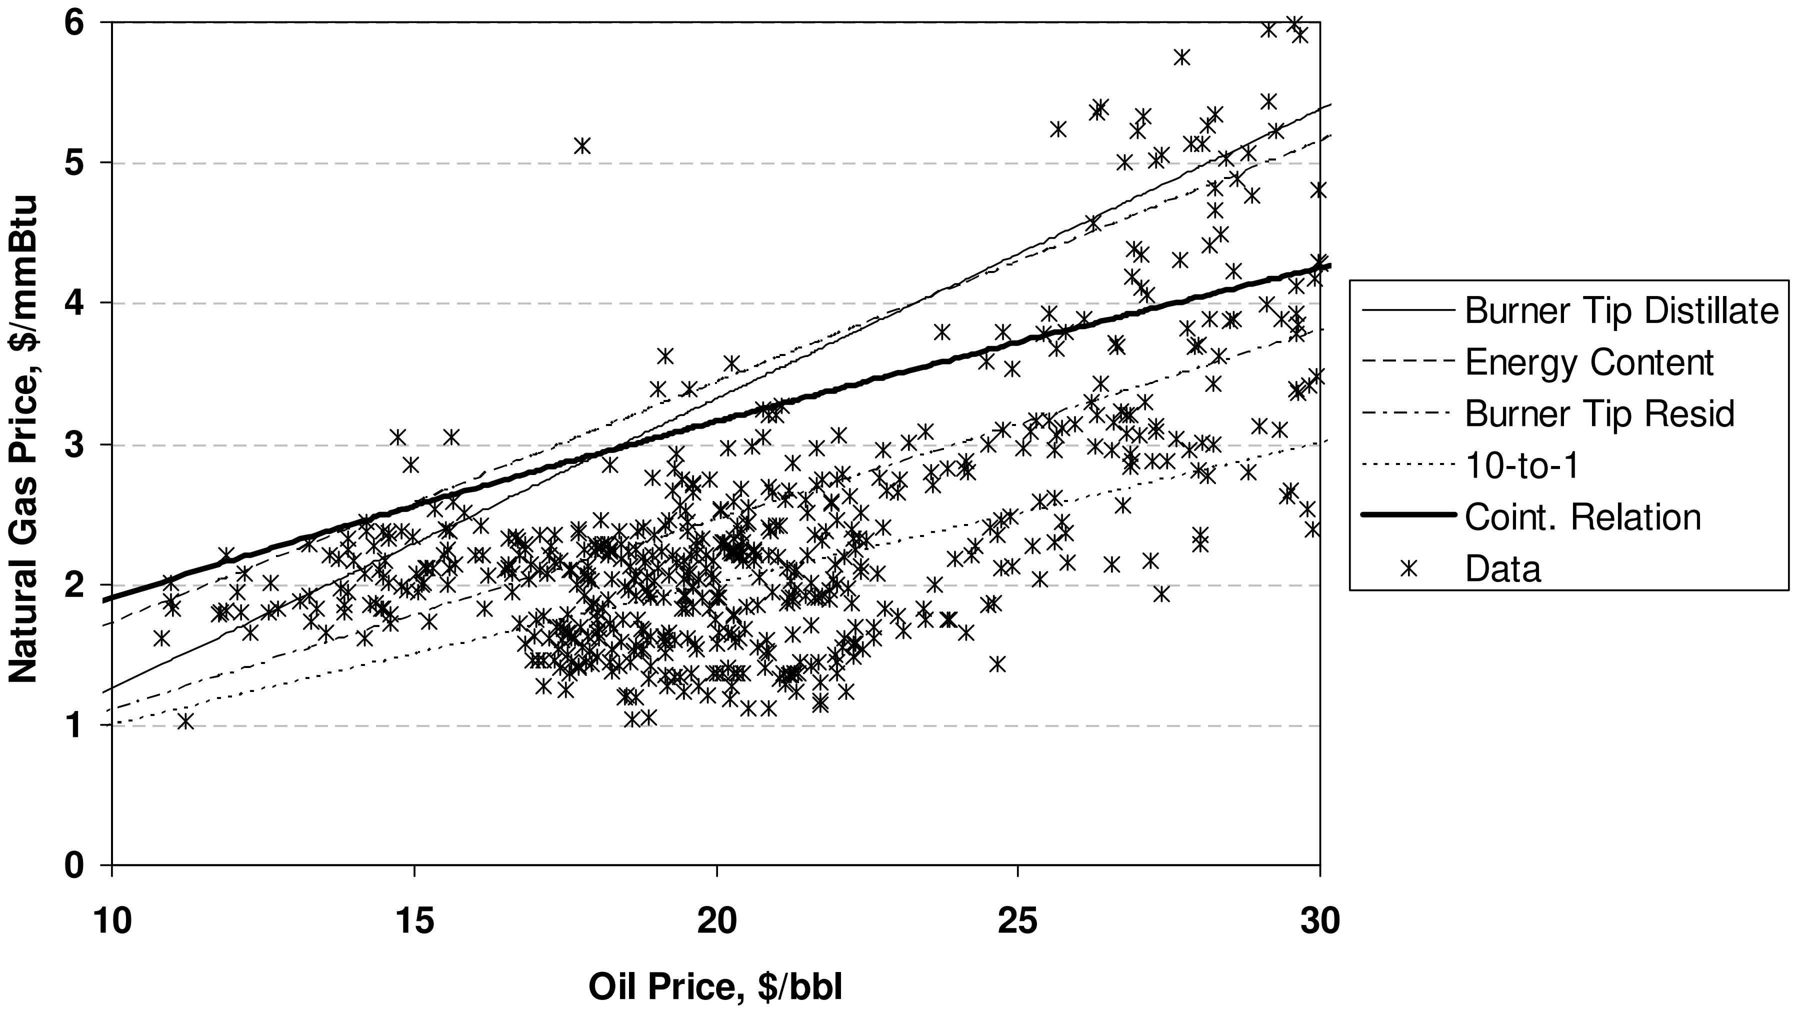

Figure 5. Prediction Errors for 4 Rules of Thumb

(Actual Log Natural Gas Price minus Predicted Log Natural Gas Price)

Fig. 5A: Burner Tip Parity Rule: Distillate

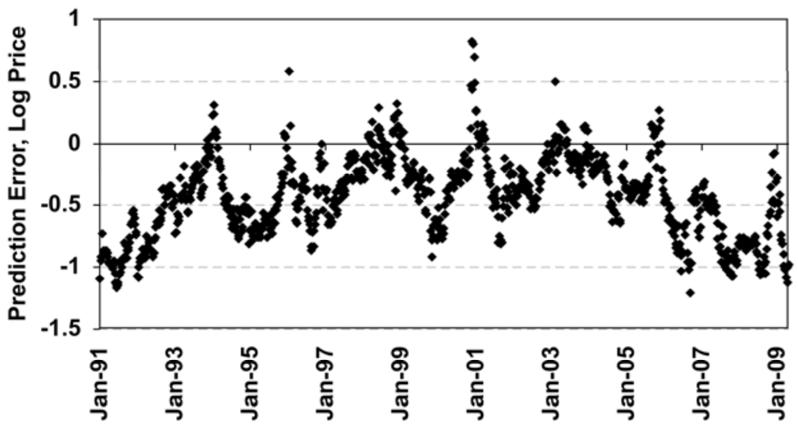

Fig. 5B: Energy Content Equivalence Rule

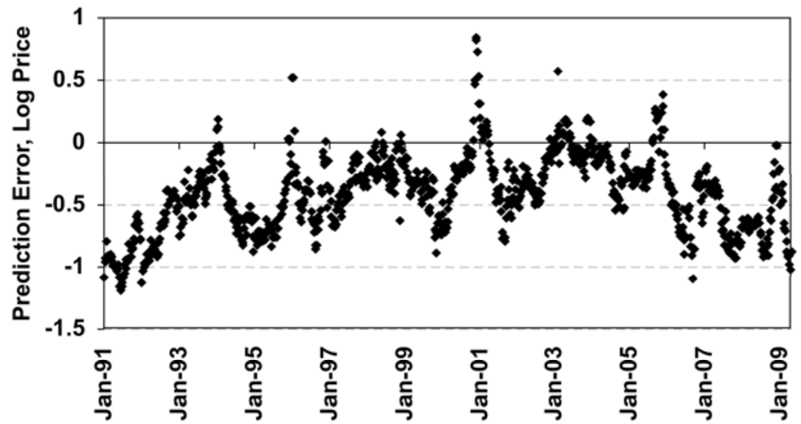

Fig. 5C: Burner Tip Parity Rule: Residual

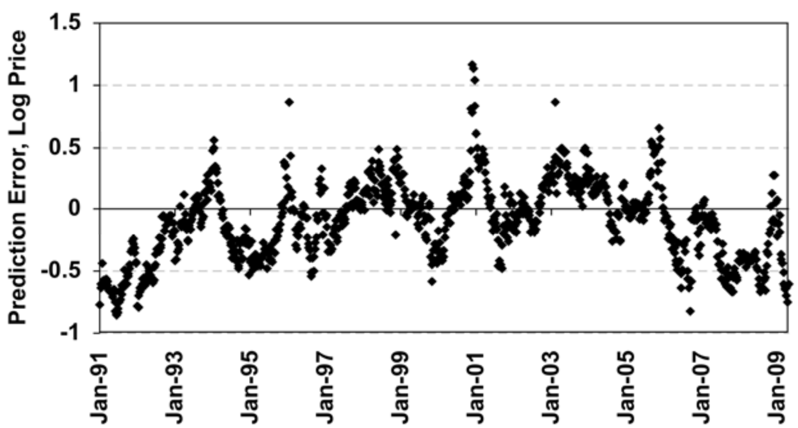

Fig. 5D: 10-to-1 Rule

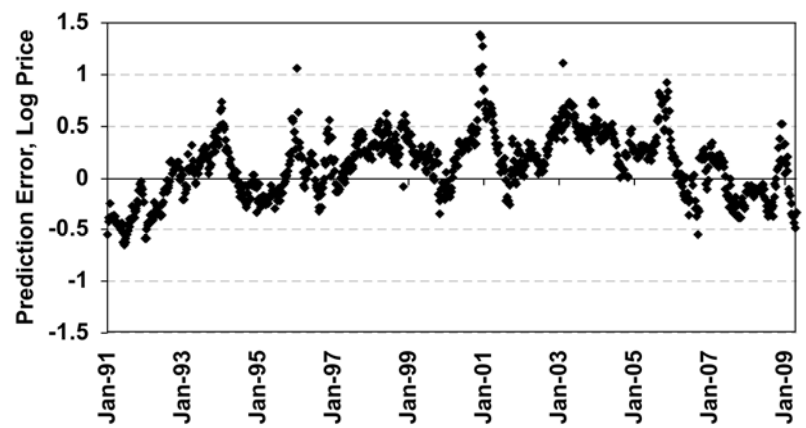

These prediction err ors

Figure 6. Prediction Errors for the Cointegrating Relationship

(Actual Log Natural Gas Price minus Predicted Log Natural Gas Price)

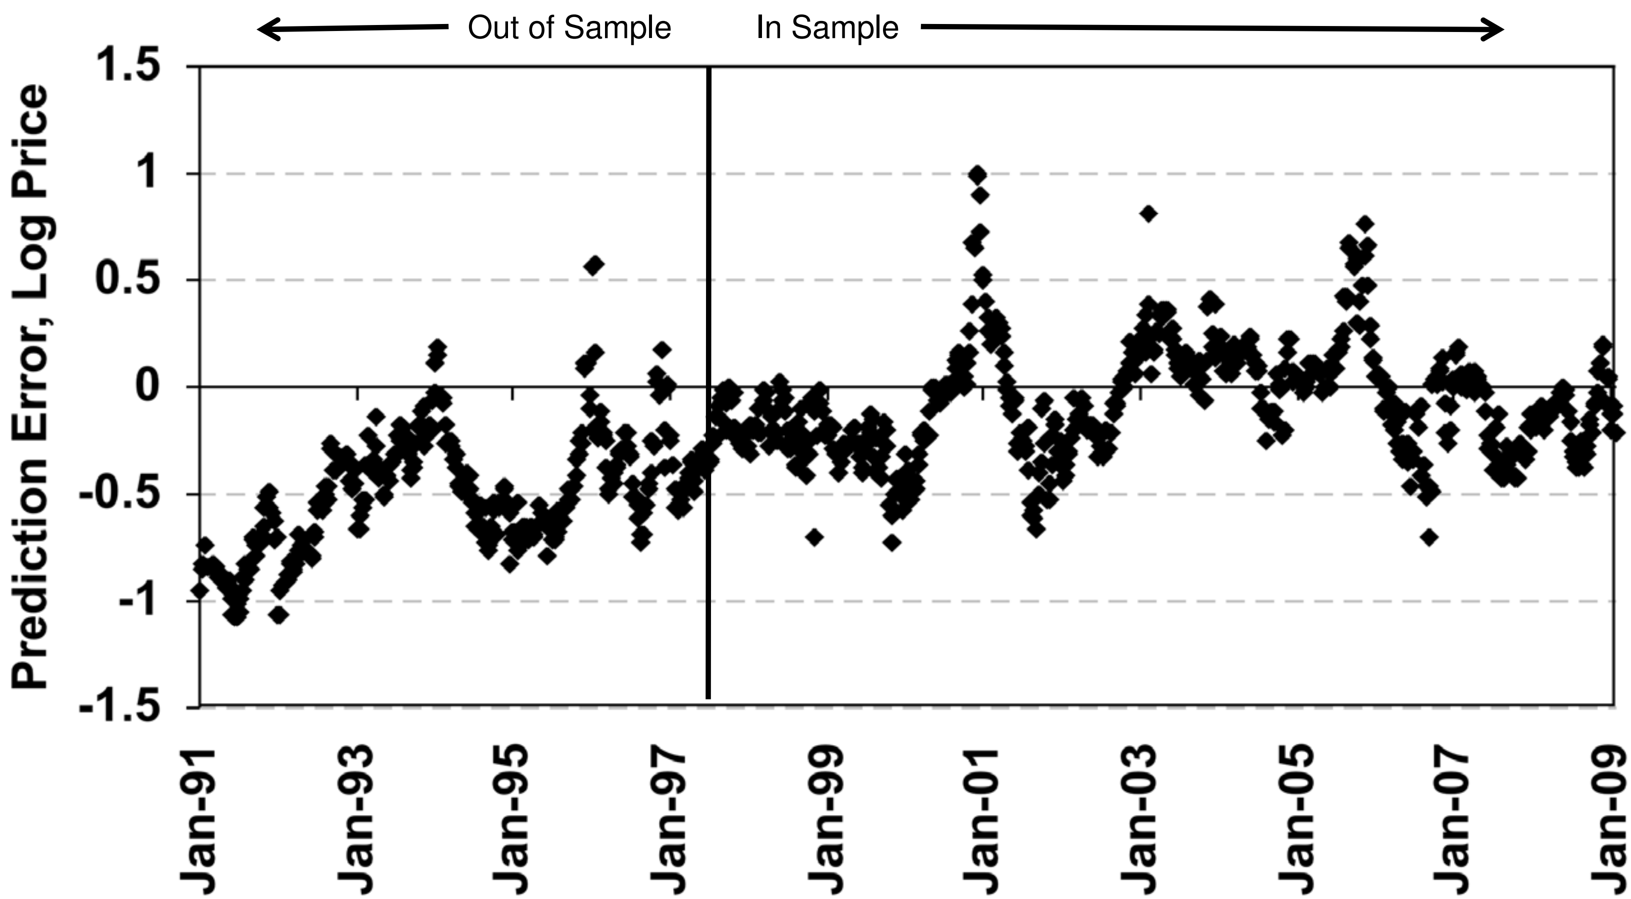

Figure 7. Seasonality in the Natural Gas Price.

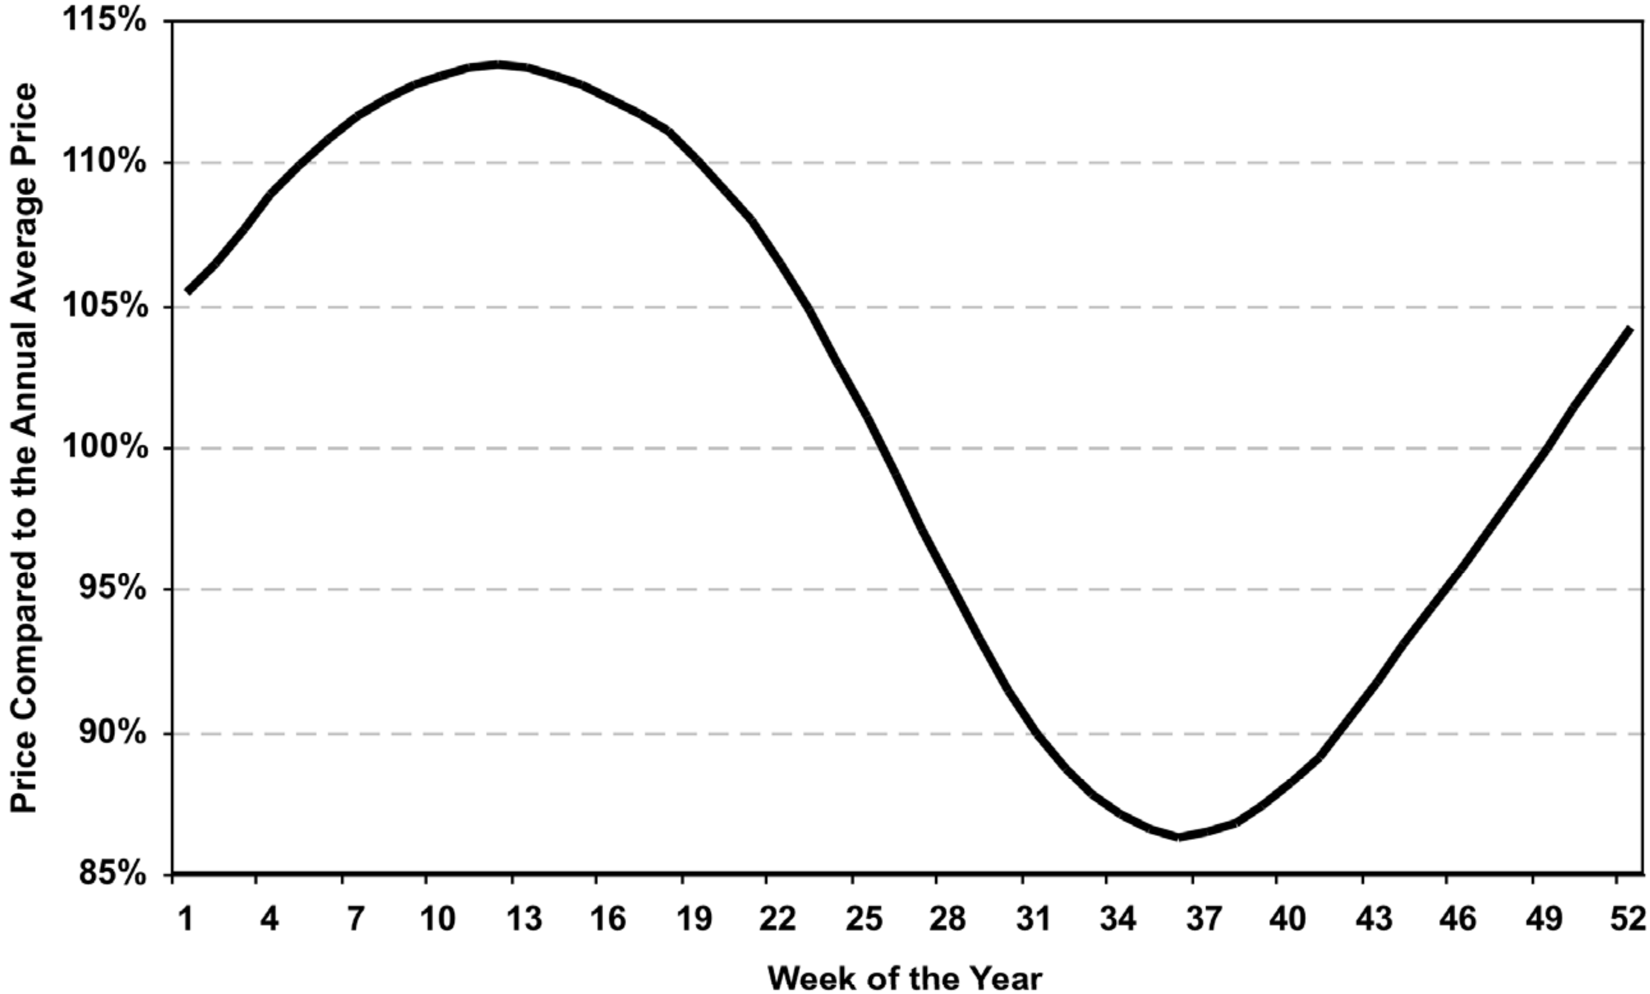

Figure 8. Seasonality Relative to Prediction Errors

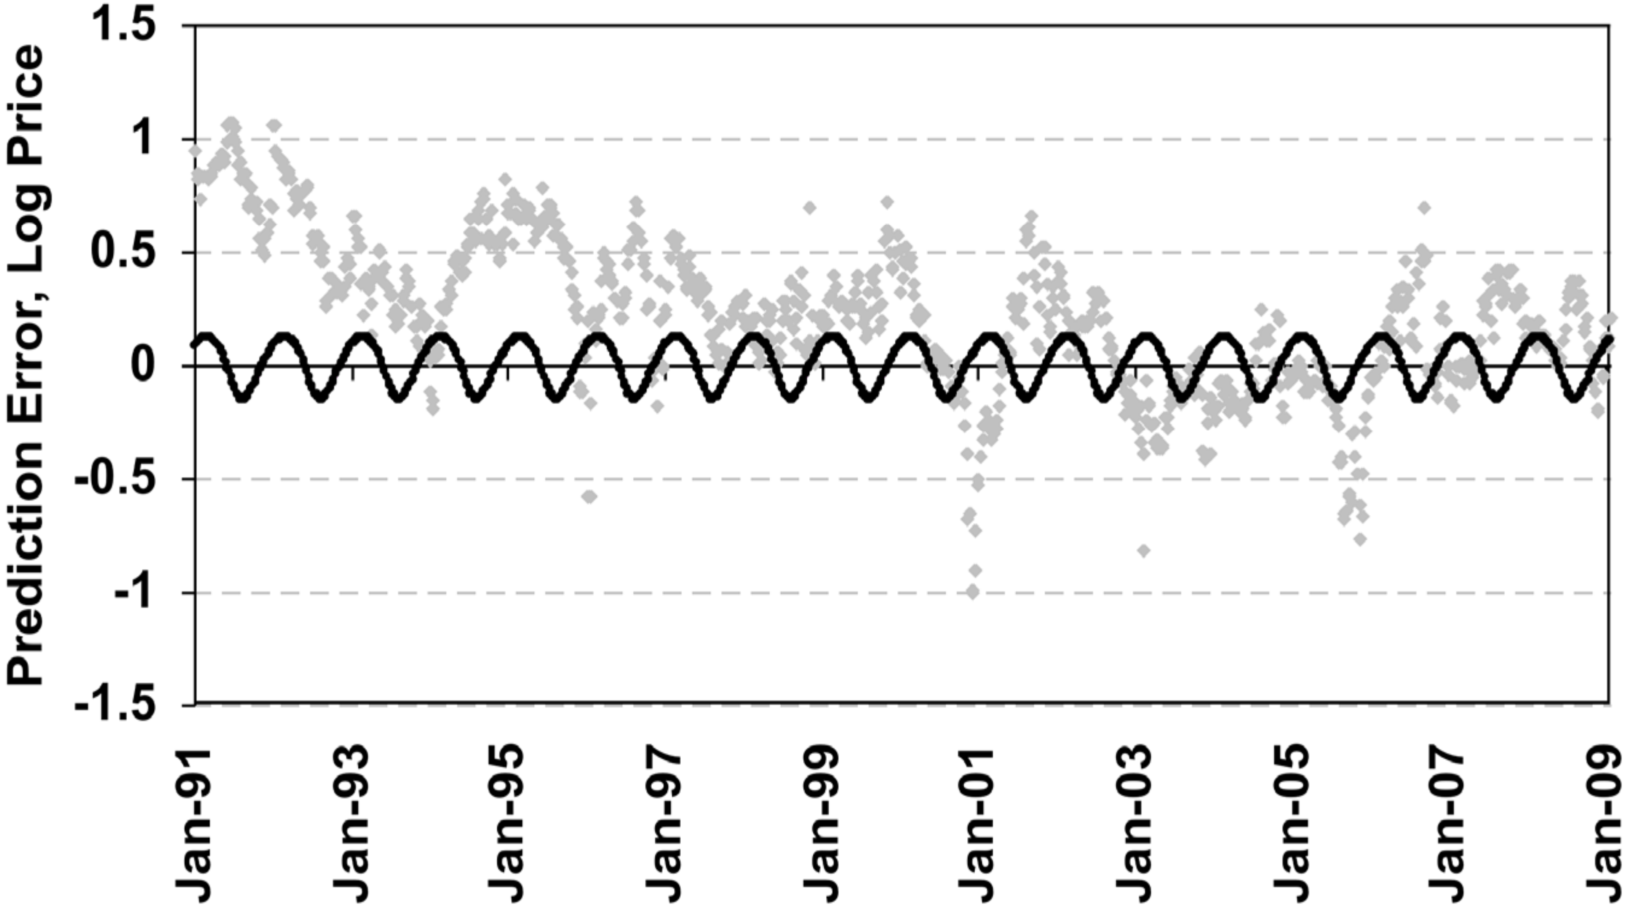

Figure 9. Impulse Response Function: The Impact of Unseasonal Weather on Natural Gas Prices.

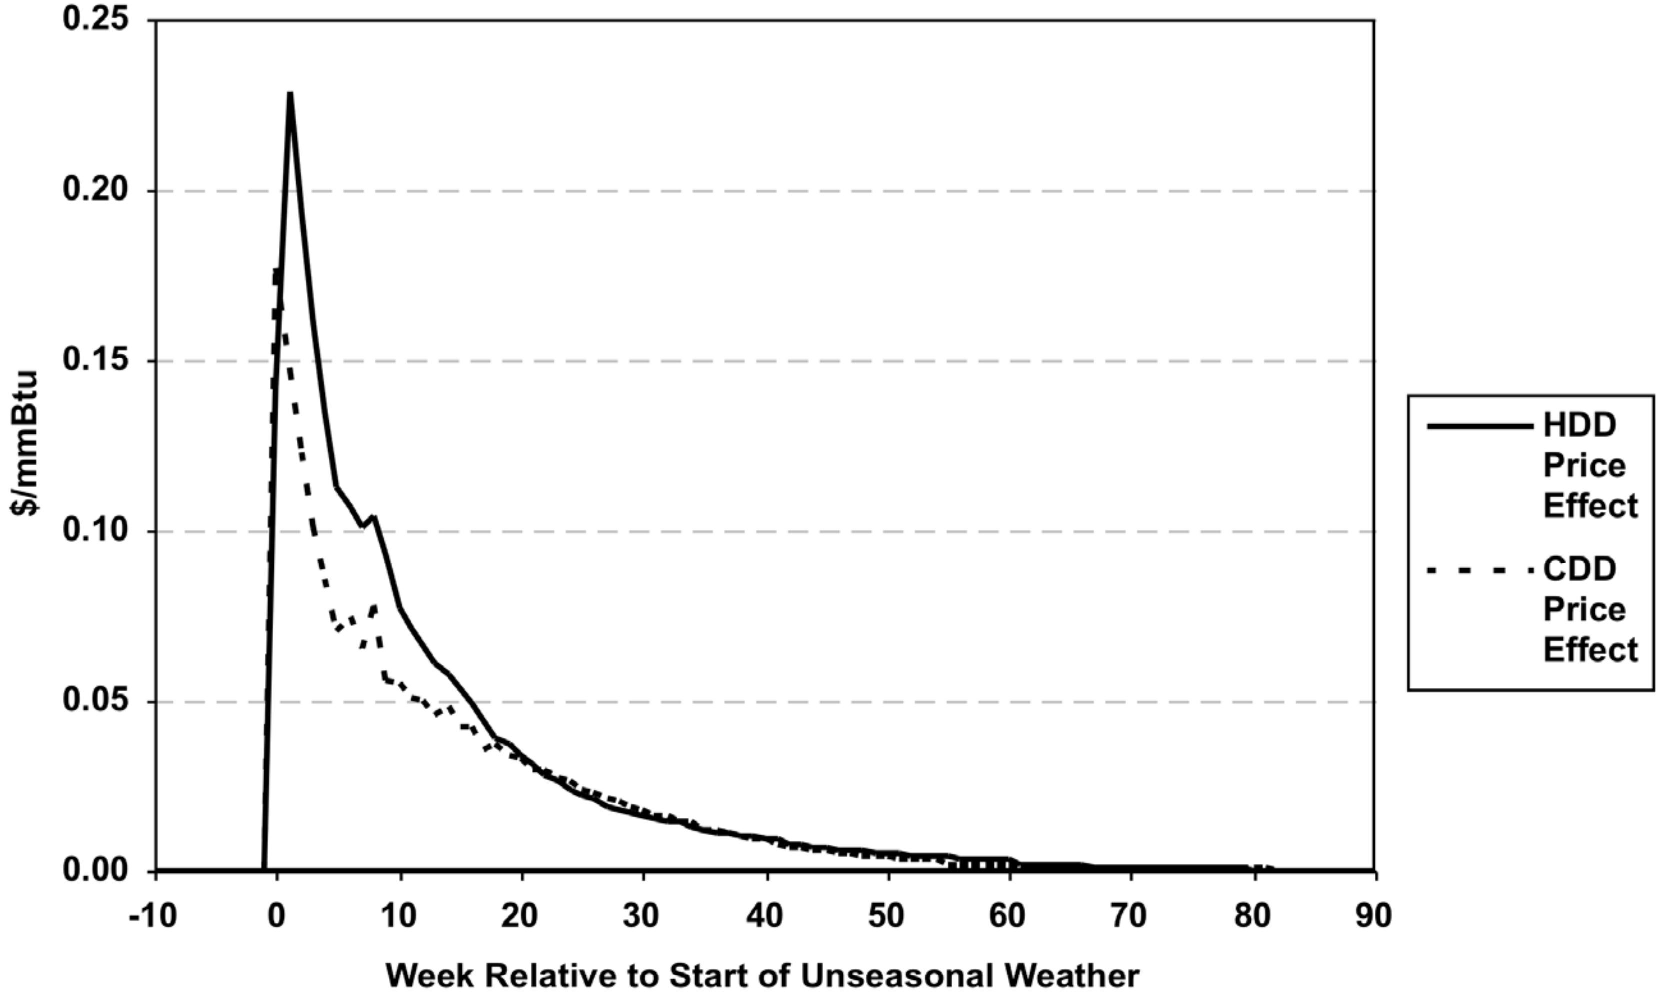

The figure shows the impulse response in the natural gas price from typical episodes of unseasonably cold winter weather (HDD) and unseasonably warm summer weather (CDD). For HDD, at week 0 the HDD is 20 degrees above normal, at week 1 the HDD is 15 degrees above normal. In all other weeks HDD is normal. For HDD, week 0 is at the start of March. For CDD, at week 0 the CDD is 10 degrees above normal. In all other weeks CDD is normal. For CDD, week 0 is at the start of August. In both cases, the impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price.

Figure 10. Impulse Response Function: The Impact of Unseasonal Cold Snap on Natural Gas Prices Against Actual Price Change.

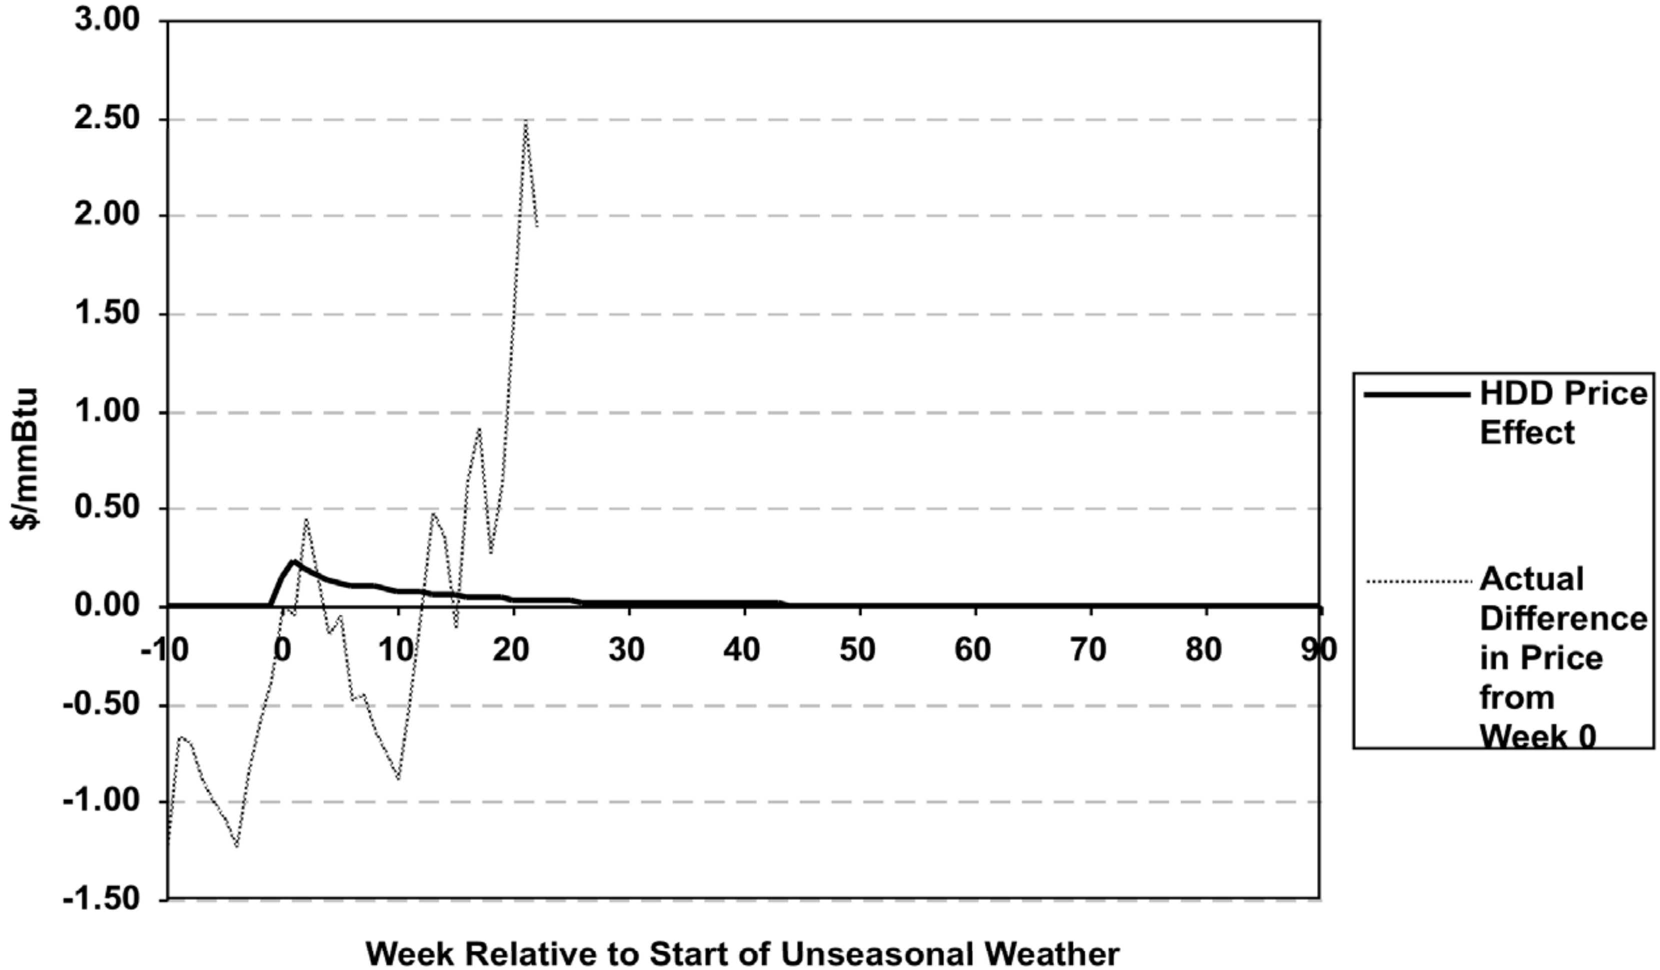

The figure shows the impulse response in the natural gas price from typical episodes of unseasonably cold winter weather (HDD). At week 0 the HDD is 20 degrees above normal, at week 1 the HDD is 15 degrees above normal. In all other weeks HDD is normal. For HDD, week 0 is at the start of March. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price.

Figure 11. Impulse Response Function: The Impact of Unseasonal Heat Wave on Natural Gas Prices Against Actual Price Change.

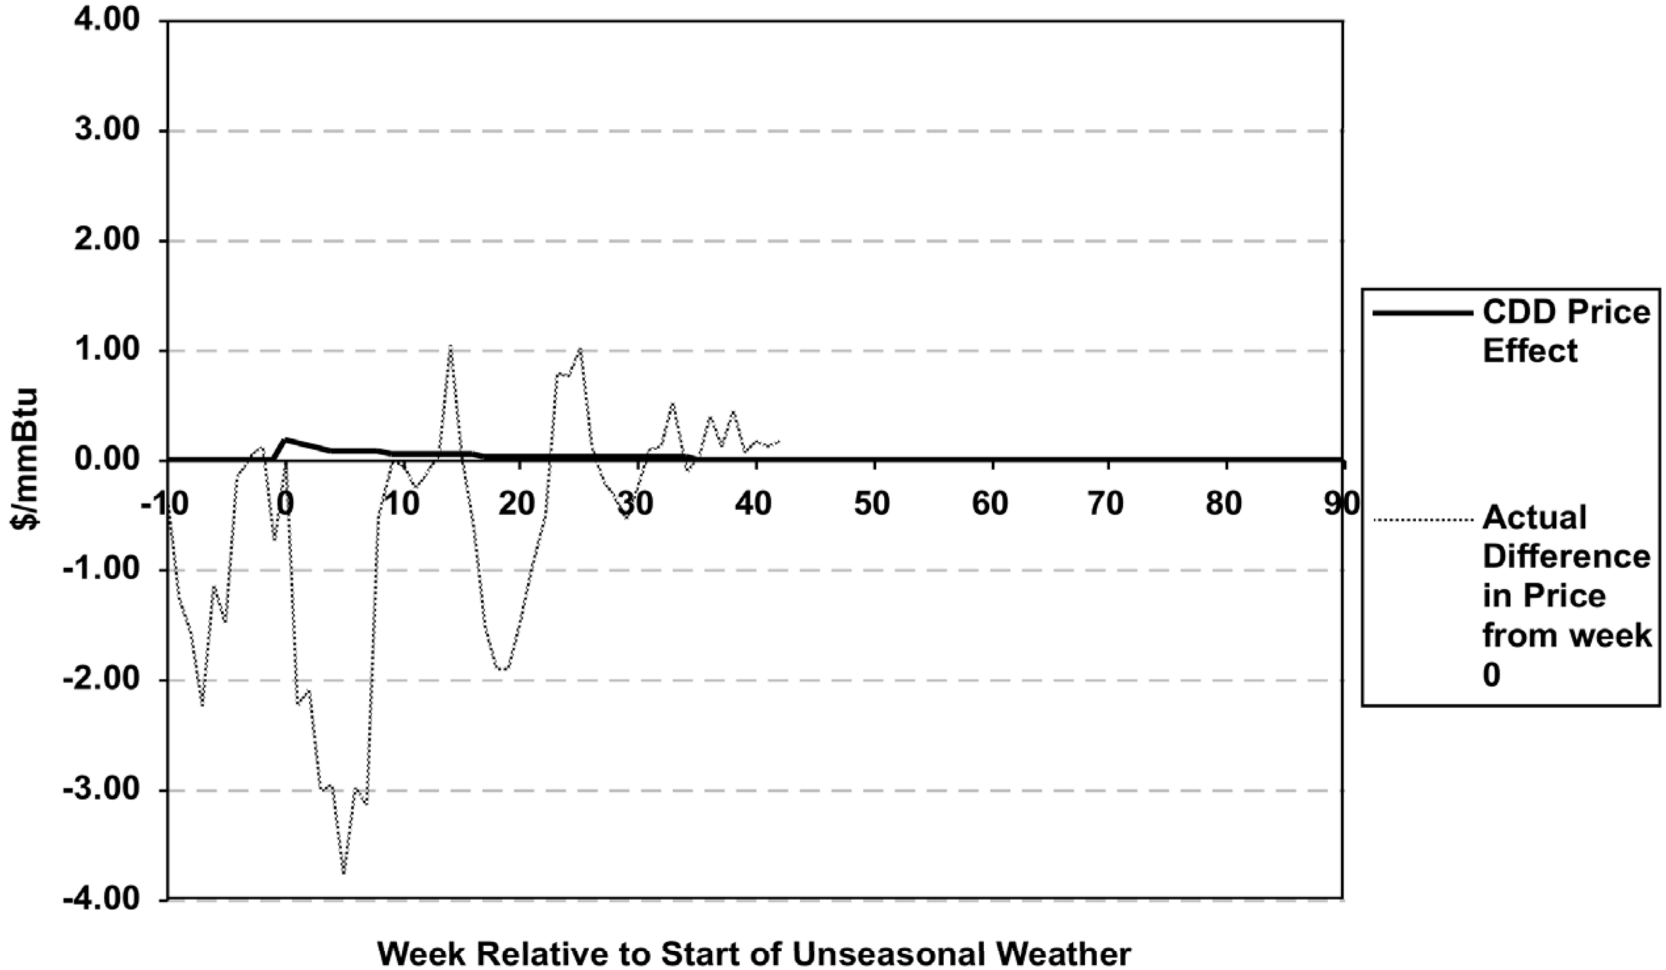

The figure shows the impulse response in the natural gas price from typical episodes of unseasonably warm summer weather (CDD). At week 0 the CDD is 10 degrees above normal. In all other weeks CDD is normal. For CDD, week 0 is at the start of August. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price.

Figure 12. Impulse Response Function: The Impact of Shut-in Production (Hurricanes).

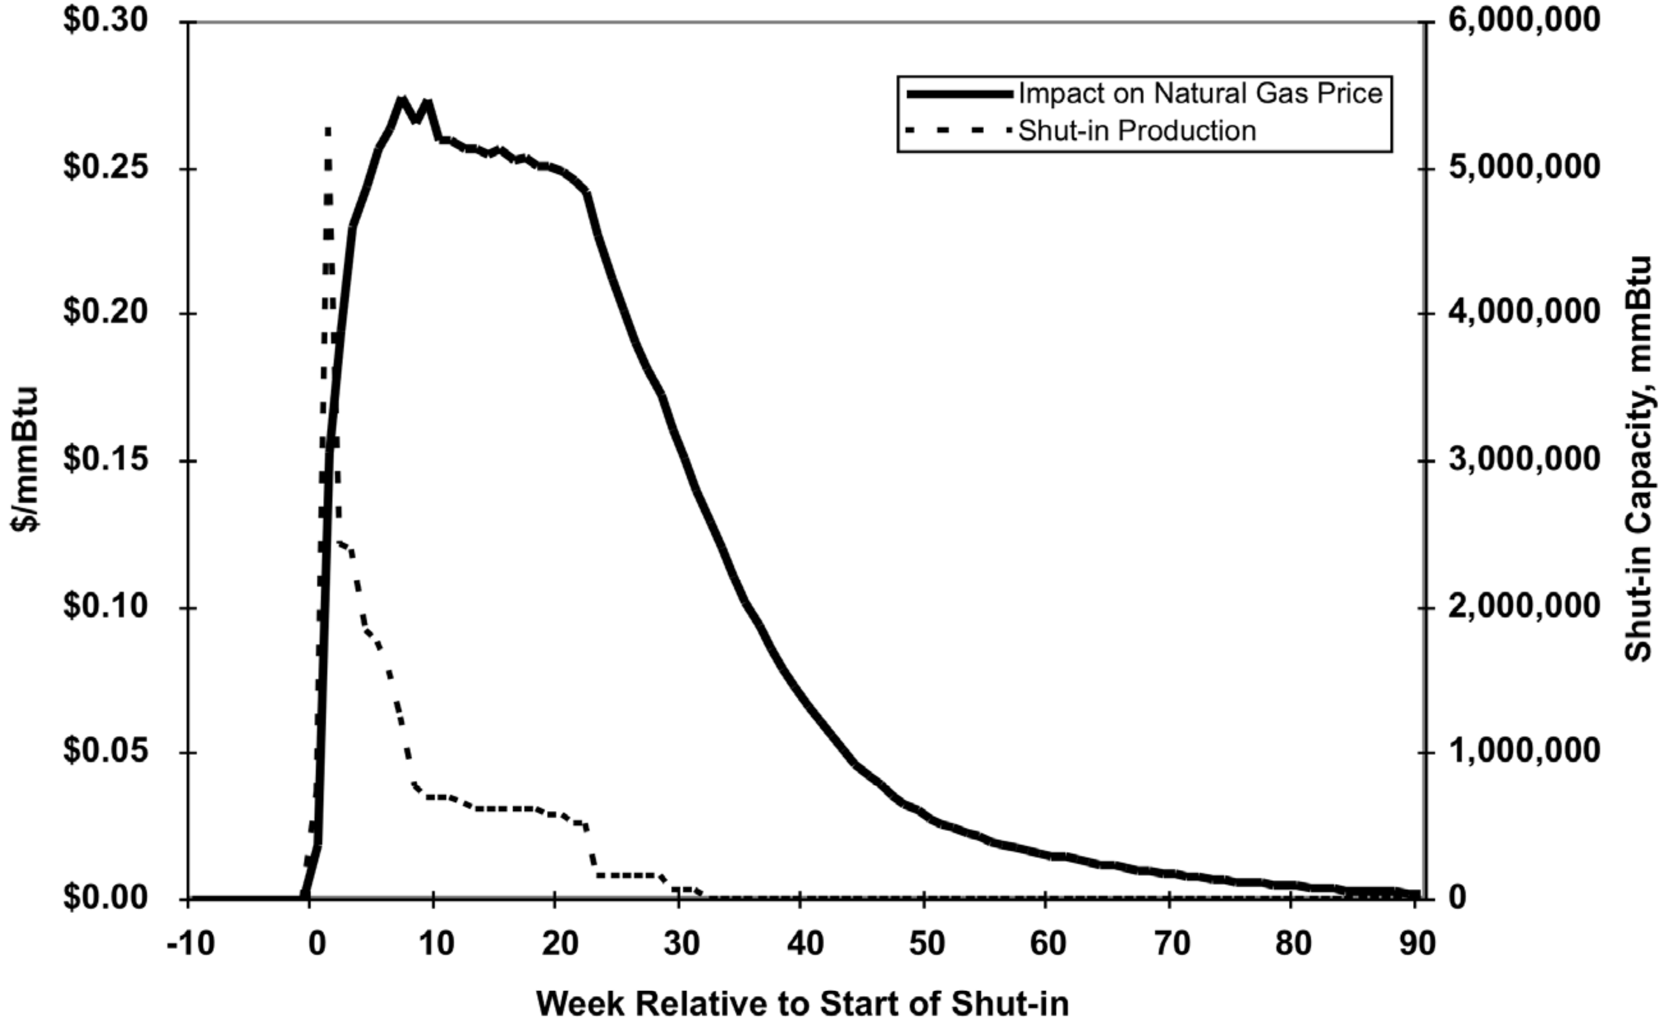

The figure shows the impulse response in the natural gas price from a period of shut-in production like that following Hurricane Ivan in 2004. At week 0 the shut-in begins. The figure shows the scale and duration of the shut-in. The maximum production shut-in is just over 5 Bcf. The shut-in lasts over 33 weeks, with an average shut-in of 0.812 Bcf. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price. Week 0 is at mid-September. The Actual Price Change series is the price difference from the actual prevailing natural gas price at week zero.

Figure 13. Impulse Response Function: The Impact of Shut-in Production (Hurricanes) Against Actual Price Change.

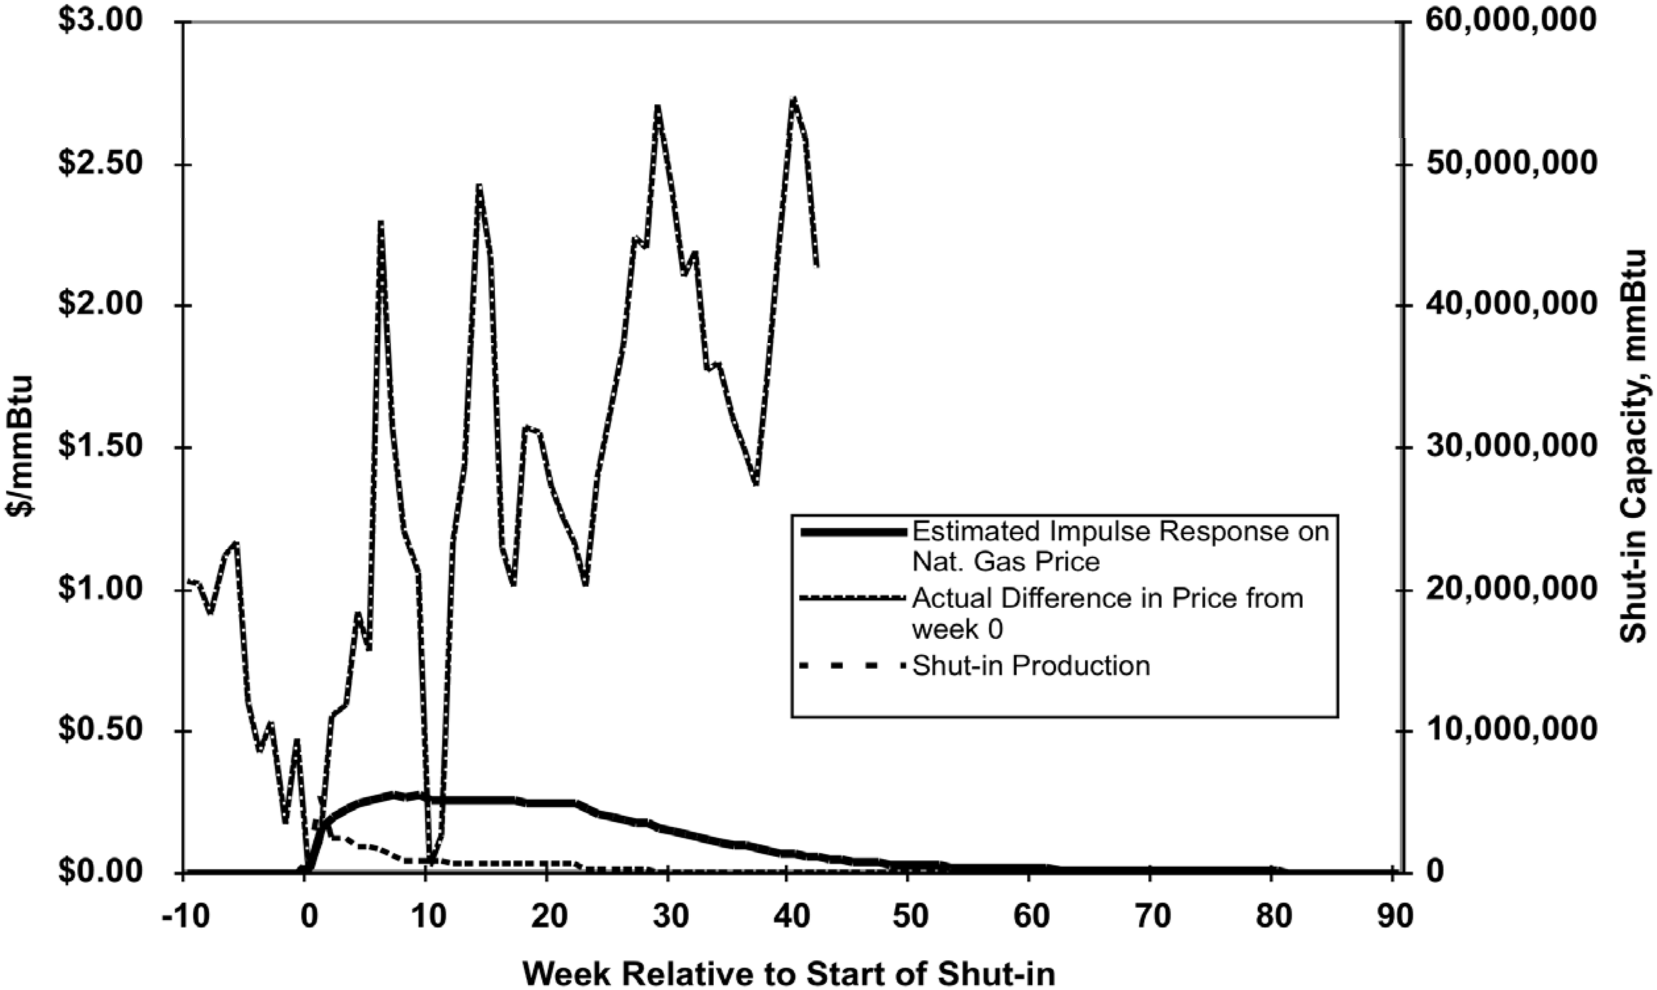

The figure shows the impulse response in the natural gas price from a period of shut-in production like that following Hurricane Ivan in 2004. At week 0 the shut-in begins. The figure shows the scale and duration of the shut-in. The maximum production shut-in is just over 5 Bcf. The shut-in lasts over 33 weeks, with an average shut-in of 0.812 Bcf. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price. Week 0 is at mid-September. The Actual Price Change series is the price difference from the actual prevailing natural gas price at week zero.

Figure 14. Impulse Response Function: The Impact of a Storage Shortfall.

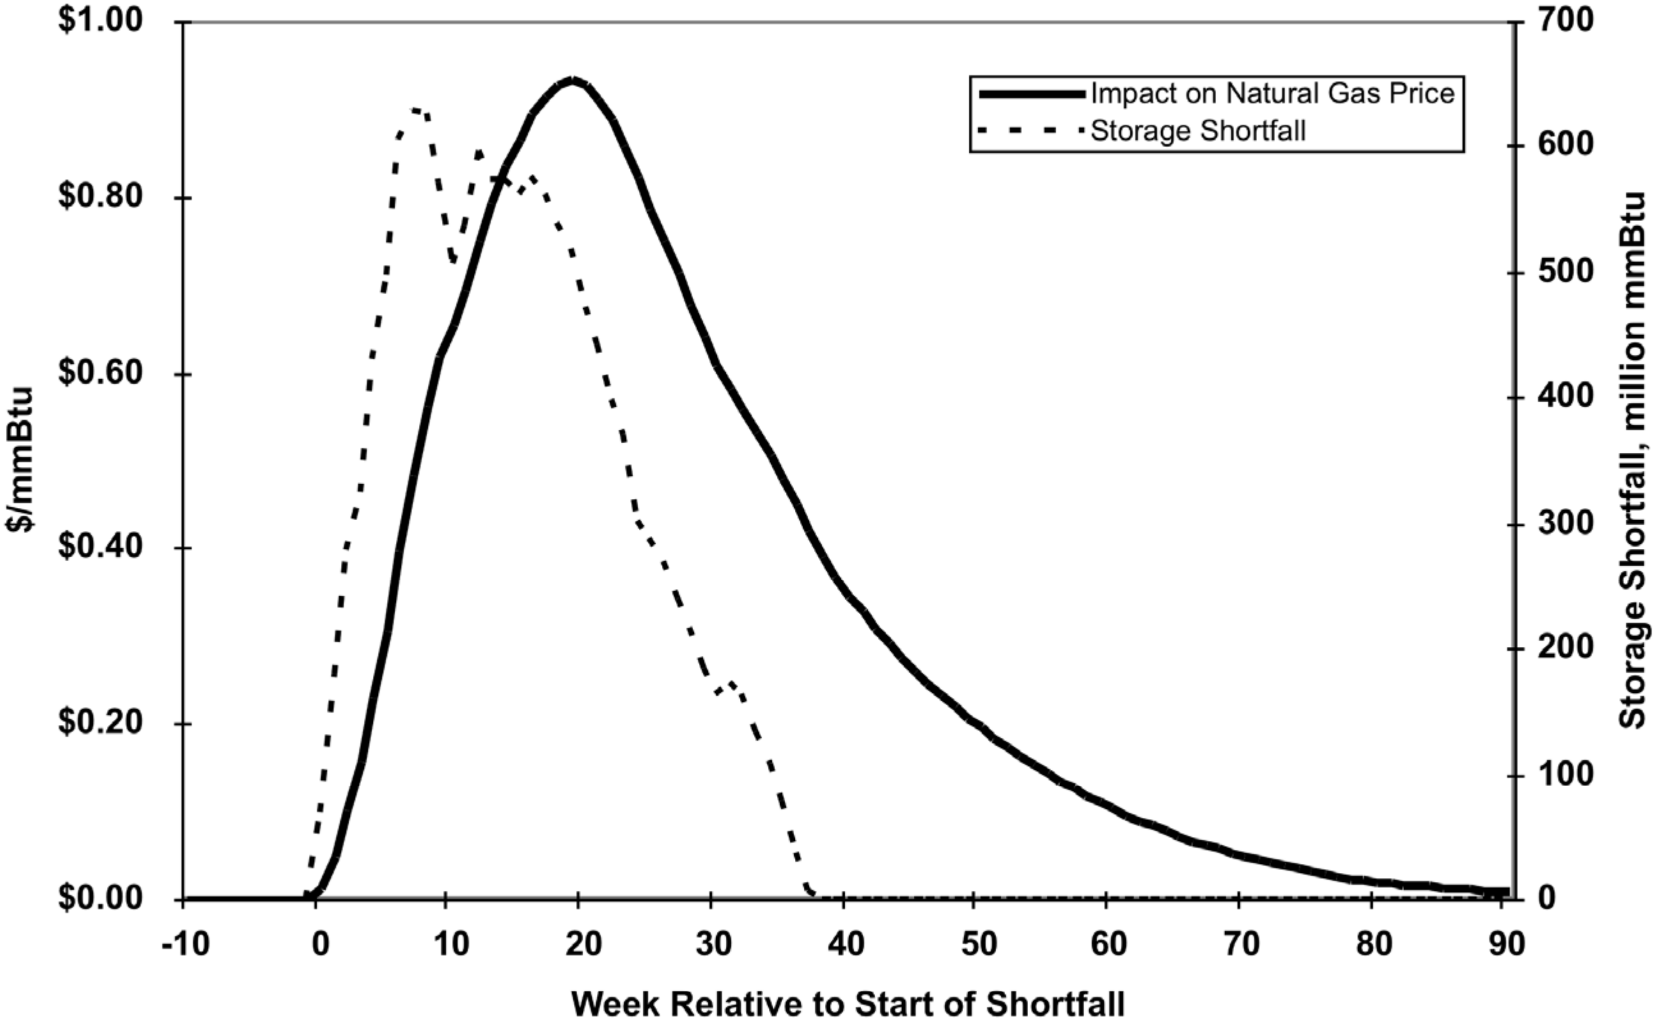

The figure shows the impulse response in the natural gas price from a period when gas in storage is less than normal for that calendar date. At week 0 the storage shortfall begins. The figure shows the scale and duration of the shortfall. The shortfall accumulates to a maximum of just over 500 Bcf, then plateaus and gradually dissipates to zero at the 37th week. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price. Week 0 is at mid-January.

Figure 15. Impulse Response Function: The Impact of a Storage Shortfall Against Actual Price Change.

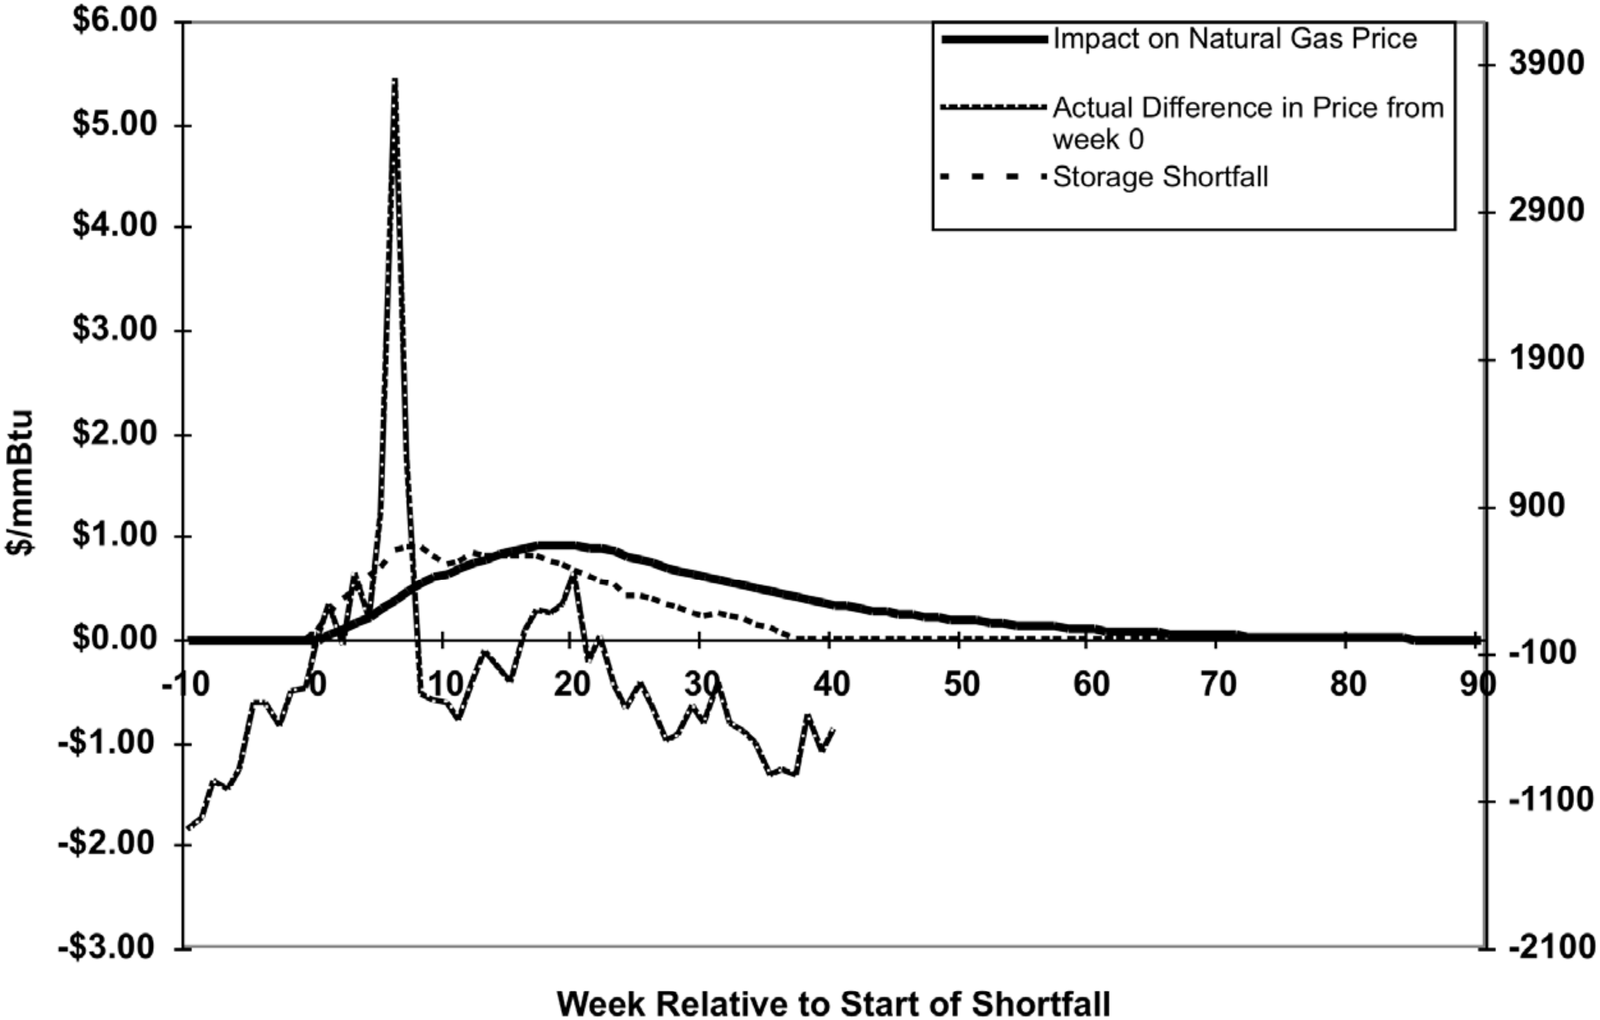

The figure shows the impulse response in the natural gas price from a period when gas in storage is less than normal for that calendar date. At week 0 the storage shortfall begins. The figure shows the scale and duration of the shortfall. The shortfall accumulates to a maximum of just over 500 Bcf, then plateaus and gradually dissipates to zero at the 37th week. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price. Week 0 is at mid-January.

Figure 16. The Response of the Natural Gas Price to an Increase in the Crude Oil Price.

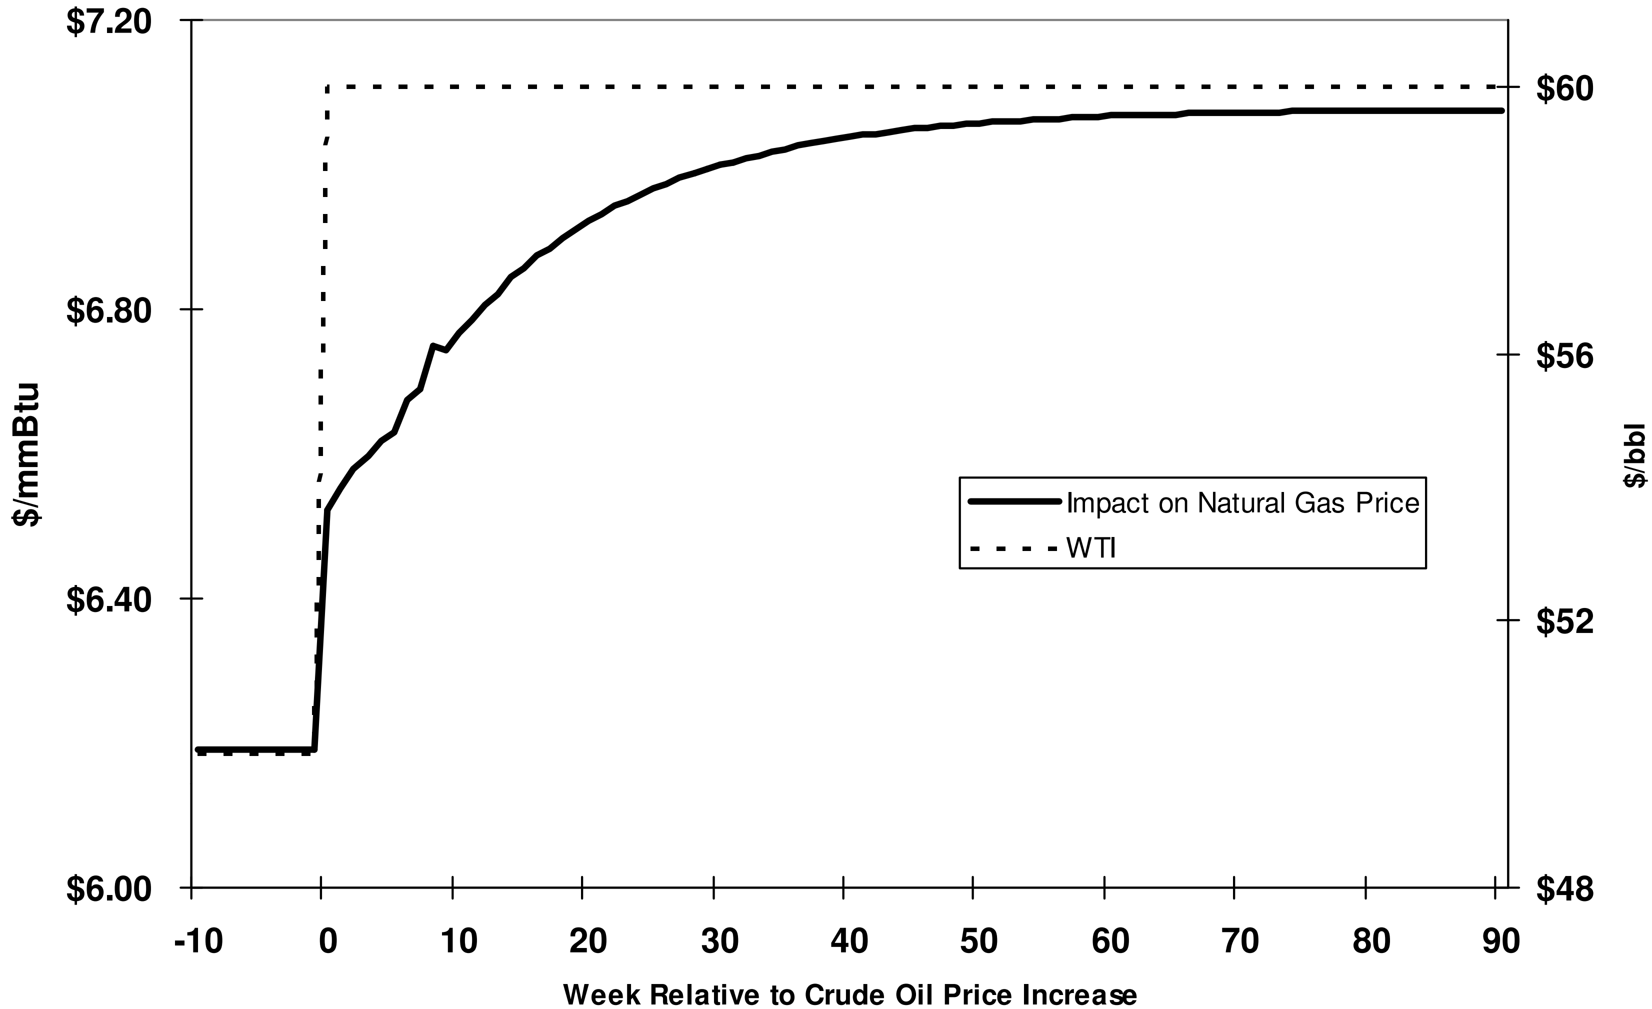

The figure shows the impulse response in the natural gas price from a $10/bbl increase in the crude oil price, from $50/bbl to $60/bbl. The impulse response is measured as compared against the seasonally adjusted natural gas price.

Figure 17. Dissipation of a Temporary Shock to the Natural Gas Price.

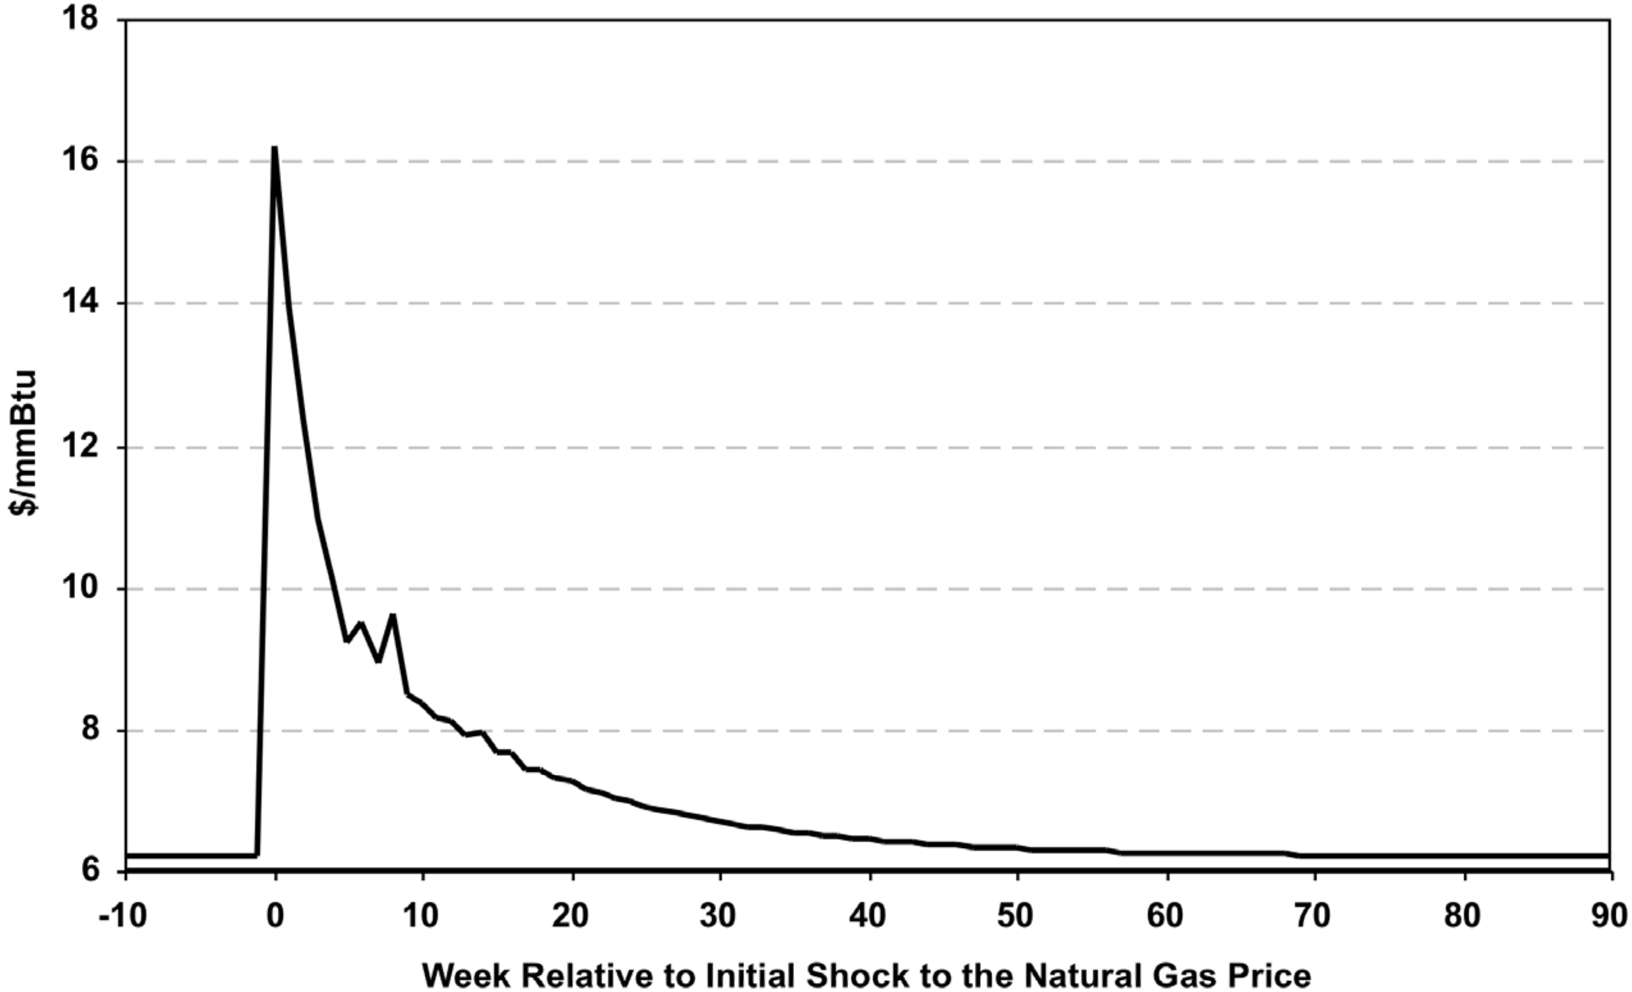

The figure shows the impulse response in the natural gas price following an exogenous initial increase of $10/mmBtu, from $6.19/mmBtu up to $16.19/mmBtu. The impulse response is shown assuming a base crude oil price of $50/bbl and therefore a base natural gas price of $6.19/mmBtu. The impulse response is measured as compared against the seasonally adjusted natural gas price.

Figure 18. Villar-Joutz Oil-Gas Cointegrating Relationship, 1989 and 2005

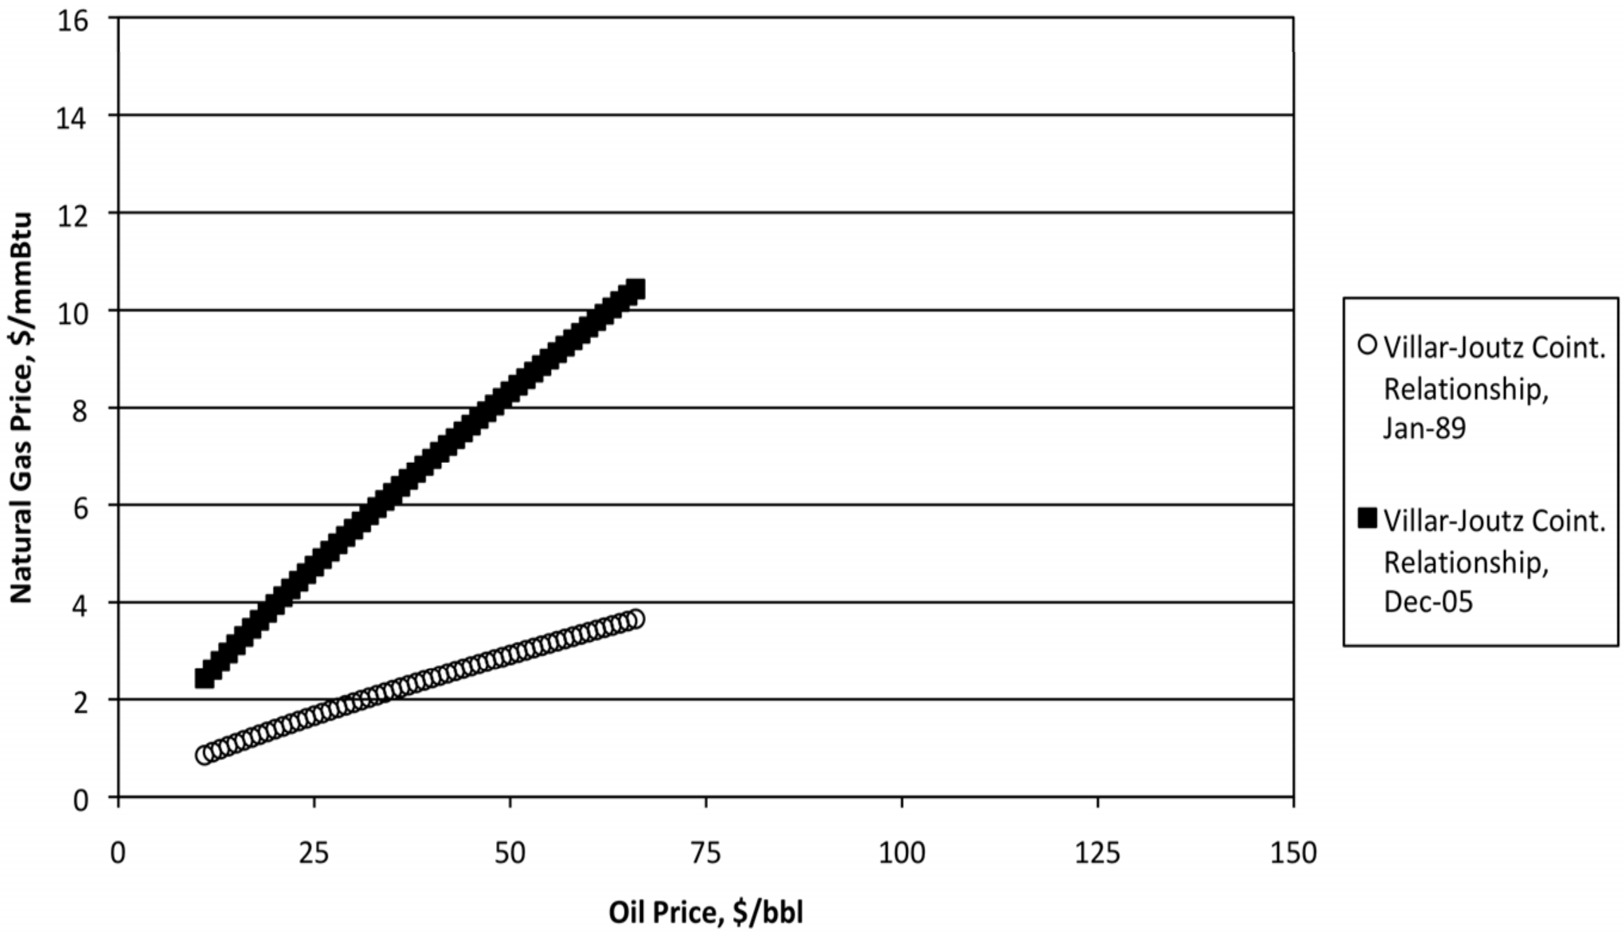

Figure 19. Segmented Oil-Gas Cointegrating Relationships (VECM)

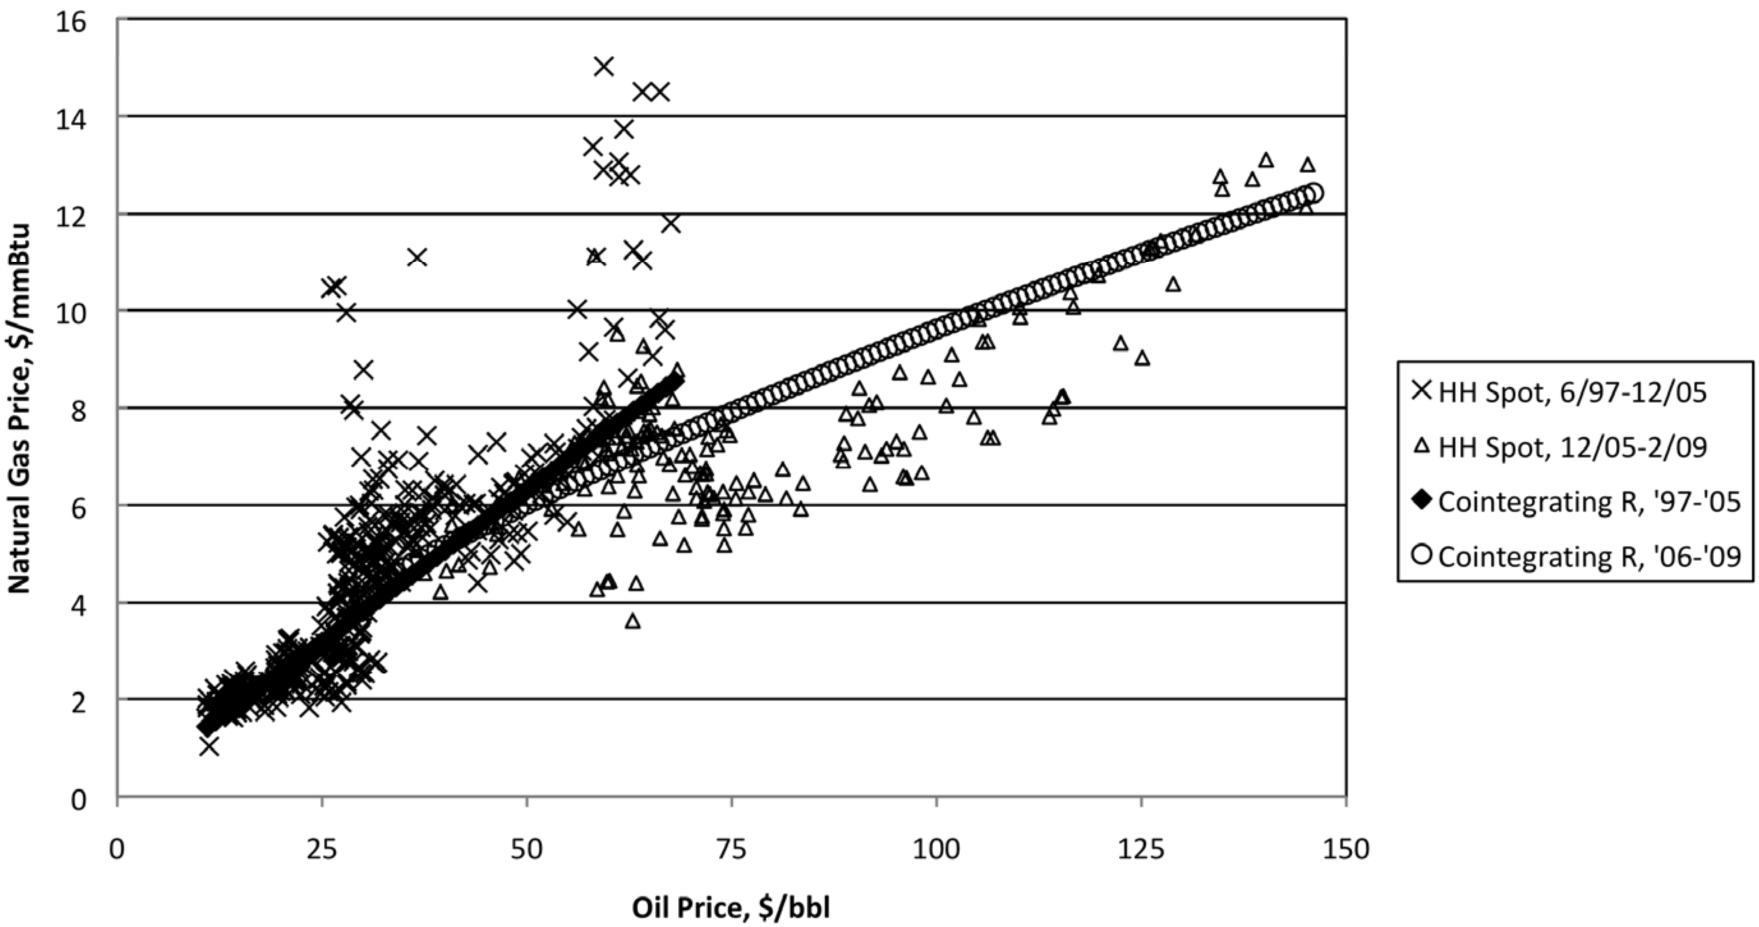

Appendix

A1. THE DATA

Our weekly price series for natural gas and for crude oil run from January 25, 1991 through February 20, 2009. 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 February 20, 2009. To analyze how the fundamental relationship may have changed over time, we broke this time period up into two pieces with one running from June 13, 1997 through December 16, 2005 and the second running from December 23, 2005 through February 20, 2009.

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 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. All of our analyses were conducted using the *natural logs* of the two price series. Any reference to the natural gas or oil price in this Appendix pertains to the natural logs of these 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]

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.

[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 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. The Augmented Dickey-Fuller coefficient for the natural log of the Henry Hub price series was -2.441, with a p-value of 13.1%. The ADF coefficient for the natural log of the WTI crude oil price was -1.477, which corresponds to a p-value of 54.5%. Identical conclusions can be drawn from the Phillips-Perron test results. For the logged Henry Hub natural gas prices, the Z(rho) statistic is -9.103 and the Z(t) statistic is -2.262. These correspond to a p-value of 18.4%. In the case of logged WTI crude oil prices, the Z(rho) statistic is -3.204 and the Z(t) statistic is -1.399. These correspond to a p-value of 58.3%. 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. 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 WTI crude oil and Henry Hub natural gas prices were stationary at the 1% level. For the Henry Hub natural gas prices, the first differences yielded an ADF test statistic of -24.984 and a Phillips-Perron Z(rho) statistic of -558.254. The Phillips-Perron z(t) statistic was -25.164. The equivalent statistics for the differenced WTI crude oil price series were -27.185, -664.208, and -27.207 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.

Seasonality – Heating Degree Days and Cooling Degree Days

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 1997-2007 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 1997-2009 data series, the HDD variable ranged from an average high of 205 in January to an average low of 1 in July. The average HDD figure was 87.8, with a standard deviation of 79.8. The CDD variable typically ranged from a low of about 1 in January to a high of about 75 in July. CDD figures were much less volatile than HDD figures. The average CDD was 25.4, with a standard deviation of 28.2.

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 1997-2009 period, deviations from normal HDD were, on average, *negative*, at -5.9, with a standard deviation of 21.5. 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 -19.

CDDDEV shocks tended to last 1-2 weeks before returning to normal. From 1997-2009, the average CDDDEV figure was 2, with a standard deviation of 6.5. Shocks tended to be highest in September, with an average value of 5.

<u>Shut-in Gas Production – Hurricanes</u>

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 1997-2009 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, reported in billion cubic feet (Bcf, equivalent to 1.03 million mmBtu). The average storage level from 1997-2009 was 2,256 Bcf (2.32 billion mmBtu) on any given week, with a standard deviation of 716 Bcf (737.84 million mmBtu). The average storage differential from 1997-2009 was 133 Bcf (137.06 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 17 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 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 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 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 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

[2] Using the "varsoc" command on Stata 8 Intercooled.

Information Criterion (AIC) test, the Schwarz 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 two 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. Running the Johansen test using nine lags provided similar results and allowed us to draw identical conclusions. The Johansen test indicated a rank of one (at significance of 1%) based on the trace statistics, and both the SBIC and HQIC likewise implied a rank of one. This is strong evidence of a single cointegrating relationship between Henry Hub natural gas and WTI crude oil prices if exogenous variables acting on the Henry Hub price are included.

[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.

A3. THE MODELS

The VECM

We then estimated our VECM.[4] The VECM is essentially a VAR model with the addition of 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 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 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 the previously-discussed exogenous variables that act solely on our hypothesized dependent variable of gas prices (HDD, CDD, HDDDEV, CDDDEV, SHUTIN, and STORDIFF). These fundamentals partially account for movements of the Henry Hub natural gas price either closer to or further from the calculated long-run target.

As stated above, the difference between the actual Henry Hub price and the long-run equilibrium 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 Henry Hub price *change* and the model's predicted Henry Hub price *change*.) An error-correction mechanism moves the natural gas price closer to its long-term equilibrium

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

relationship with WTI 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:

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

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

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

The top equation details the long-term relationship between Henry Hub and WTI prices. $P_{HH,t}$ is the logged Henry Hub natural gas price in week t, $P_{WTI,t}$ is the logged 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 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 cointegration 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 – 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 assumption 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 illustrates the results for the VECM models assuming Henry Hub prices and, alternately, WTI prices as the dependent variable. For the single cointegrating relationship, the β coefficient is 0.7342, and the γ coefficient is -1.0493. The β coefficient is highly statistically significant, with a p-value of 0.000.

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

$$log(P_{HH,t}) = -1.0493 + 0.7342 log(P_{WTI,t})$$

The equation for changes in the Henry Hub price has an R^2 statistic of 0.1411, while the equation for changes in the WTI price has an R^2 statistic of 0.1143. This implies that in the case of Henry Hub prices, about 14.1% 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 11.4% 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 -0.0828, with a p-value of 0.0000. 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 8.4 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 STORDIFF 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 51.15. 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, only the Henry Hub price change from three weeks prior and from nine weeks prior were statistically significant within the 5% p-value range, and the second, fifth, and eighth lags were significant at the 10% level. However, jointly the nine lagged price changes have a p-value of 0.0041, which is well within the 1% significance range.

None of the lagged WTI price changes were statistically significant, either alone or jointly. The joint significance test of the WTI lagged price changes returned a p-value of 0.6333, which implies that the probability that the combined effects of the nine lagged WTI prices are actually zero are around 63%. However, when included with the entire set of all variables in the regression, the Chi^2 statistic is 75.86, 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 ECM 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 the 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 WTI price imparts an immediate effect on the logged 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 Henry Hub price minus the long-run predicted logged 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:

$$\Delta P_{HH,t} = a + \alpha(\mu_{t-1}) + b\Delta P_{WTI,t} + \sum_{i=1}^{n} c_i \Delta P_{HH,t-i} + \sum_{j=1}^{6} d_j X_{j,t} + \varepsilon_t$$

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 used 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 0.1606, meaning that over 16% 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 114.8 from its VECM value of 94.3.

The α coefficient for the error correction has intensified to -0.0952 and has remained statistically significant, with a p-value of 0.0000. The half-life for the error correction mechanism to eliminate differences from the long-run predicted relationship is 7.3 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 for each of the variables that were significant in the VECM, with the STORDIFF p-value of 0.037 bringing it into the 5% significance range and the CDD p-value improving to within the 10% significance range at 0.069. 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 0.2872, but also statistically significant at 0.0000.

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 discrepancies in the predicted price changes in Henry Hub are rather small, especially considering the mismatch in the time periods being studied and 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).

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 price of Henry Hub and the contemporaneous change in the logged price of WTI. All are used to determine how the logged 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 price at Henry Hub in $/mmBtu. The table provides an example of how a one-unit increase in the relevant variable will change the 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 Henry Hub prices or the change in the 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. The general magnitude of effect of a one-standard deviation increase on the price of Henry Hub natural gas is similar across all but three of the variables, ranging from about 3 to 6 cents/mmBtu. The three big outliers are the change in the WTI price, HDDDEV and CDDDEV. 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 Henry Hub price begins at $7/mmBtu. A one-standard deviation increase in WTI prices under the same initial conditions would provoke a 12-cent/mmBtu increase in the Henry Hub price. Clearly, the change in WTI prices is economically on par with HDDDEV and CDDDEV in terms of effect on the price at Henry Hub.

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 1997-2009 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.

A6. SEGMENTED TIME PERIOD MODELS

The methodology for creating the VECM and conditional ECM for the 1997-2005 and the 2006-2009 segments mirrors that of the full 1997-2009 period model. We will not go into too much detail here. However, we will present the relevant data and results in the following paragraphs.

The results of the Augmented Dickey-Fuller (ADF) Tests and the Phillips-Perron Tests for unit roots resulted in nearly identical conclusions for the 1997-2005 period, but this was not the case for the 2006-2009 segment. In the 1997-2005 period, only STORDIFF failed to reject the null hypothesis of a unit root outside of the 5% level, with a p-value of 11.8% for the ADF test and 13% for the Phillips-Perron test. Figures A9 and A10 present the results of the ADF and Phillips-Perron tests for unit roots over the June 13, 1997 through December 16, 2005 period.

Over the 2006-2009 period, three of the exogenous variables failed to reject the null hypothesis of a unit root at any credible level. The HDD series had a p-value of 22% in the ADF test and 21% in the Phillips-Perron test. The CDD series had a p-value of 46% under the ADF test and 22% under the Phillips-Perron test methodologies. Since both HDD and CDD are weather-related data, there is no theoretical reason that one would suspect these to have a meaningful correlation with past observations despite suggestions to the contrary by statistical inference. They were kept in the model without modification.

The STORDIFF series from December 23, 2005 through February 20, 2009 was a bit more problematic. The STORDIFF p-value in the ADF test was 58.6%, while the p-value in the Phillips-Perron test was 36.9%. Both are strong evidence of autocorrelation. However, the STORDIFF variable was retained in the model in order to provide a model as similar as possible to the models reflecting the 1997-2009 and 1997-2005 periods for purposes of comparison. Tables A11 and A12 present the results of the ADF and Phillips-Perron tests on the December 23, 2005 through February 20, 2009 period.

Another point at which the results of the pre-modeling tests diverged between the long 1997-2009 period and its shorter segments was the VAR selection order criteria tests. For the 1997-2005 period, the test suggested either seven lags or one lag. Since a VECM requires at least two lags, and since more tests selected seven lags, we used seven lags in our segmented model. Table A13 presents the results of the VAR selection order criteria test for the 1997-2005 segment.

The VAR selection order criteria tests provided less decisive results for the 2006-2009 segment. Two of the tests selected four lags, one test selected ten lags, one test selected three lags and one test selected a single lag. We used four lags, since it had a plurality of test results supporting it. Four lags also had stronger statistical significance when conducting the Johansen test (discussed below) for cointegration than three lags or one lag. Table A14 presents the data from the VAR selection order criteria test for the December 23, 2005 through February 20, 2009 period.

Both segments of the longer 1997-2009 period were subjected to the Johansen cointegration tests to determine if there was a statistically significant, stable relationship between crude oil and natural gas prices. The test results showed cointegration between the crude oil and natural gas prices at the selected lag lengths discussed above. The results are detailed in Table A15 for the 1997-2005 segment and in Table A16 for the 2006-2009 segment.

The coefficients and results of the VECM for both segments are detailed in Table A17 for the 1997-2005 period and in Table A18 for the 2006-2009 period. Each table also includes the joint significance of various groupings of coefficients, as well as the overall fit and statistical significance-related statistics for each model. For the 1997-2005 segment, the R^2 improved from about 14% to over 18%, while the Chi^2 statistics remained highly statistically significant. The results for the 2006-2009 segment were slightly less clear-cut, but the R^2 improved to nearly 30%, the root mean squared error improved by over one percentage point and the grouped coefficients all continued to exhibit statistical significance within the 10% level. This is despite the considerably smaller number of observations with which to conduct statistical testing.

The coefficients and results of the conditional ECM for both segments are detailed in Table A19 for the 1997-2005 period and in Table A20 for the 2006-2009 period. The conditional ECM, in both cases, has a more statistically significant fit with the data than the VECM model for each segment. Furthermore, the R^2, Chi^2, and joint significance test statistics for the 1997-2005 segment all show either an improvement over or parity with the corresponding statistics for the 1997-2009 period. The 2006-2009 period does not achieve higher statistical significance in every category against the longer model. This is undoubtedly at least partially due to the fact that there are such fewer data points with which to analyze. However, the joint significance tests are all at the same significance level and the R^2 for the segment is the strongest of all three models at nearly 31%. The modeling exercise appears to perform best both at marking the cointegrating relationship and accounting for the idiosyncratic volatility in natural gas prices in this 2006-2009 segment.

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

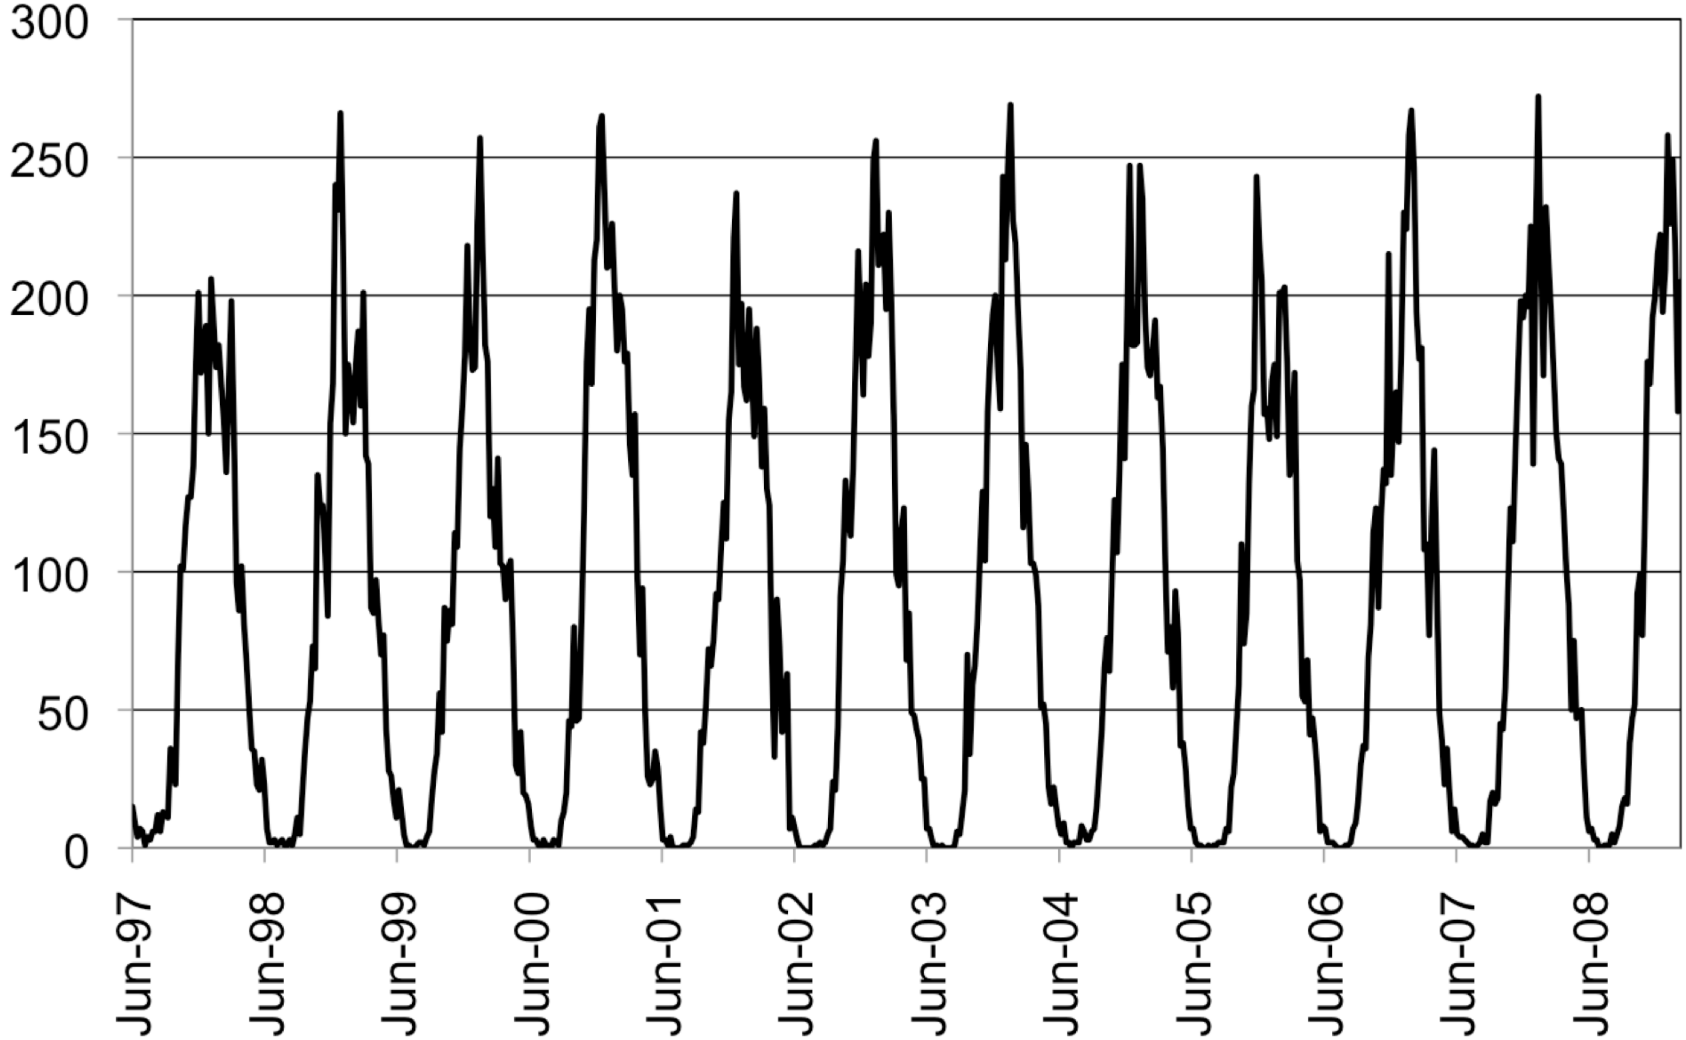

Figure A2. US Average Weekly Cooling Degree Days (CDD)

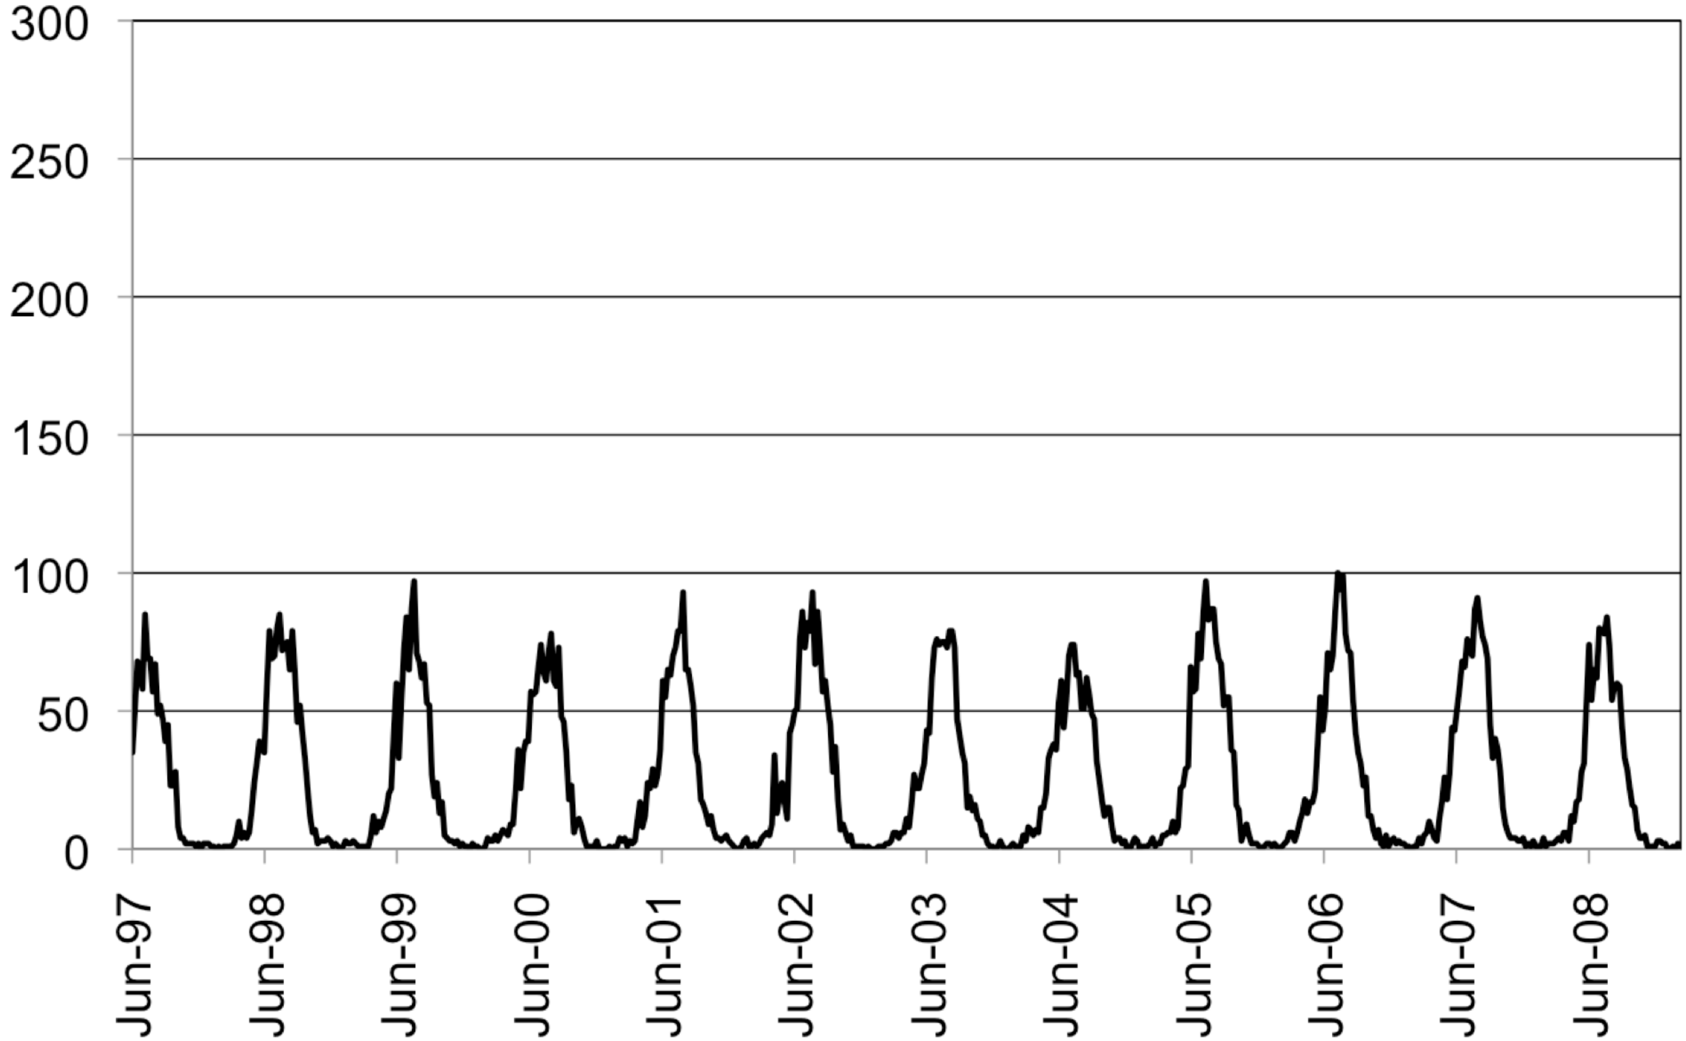

Figure A3. US Average Deviation from Normal Heating Degree Days (HDDDEV)

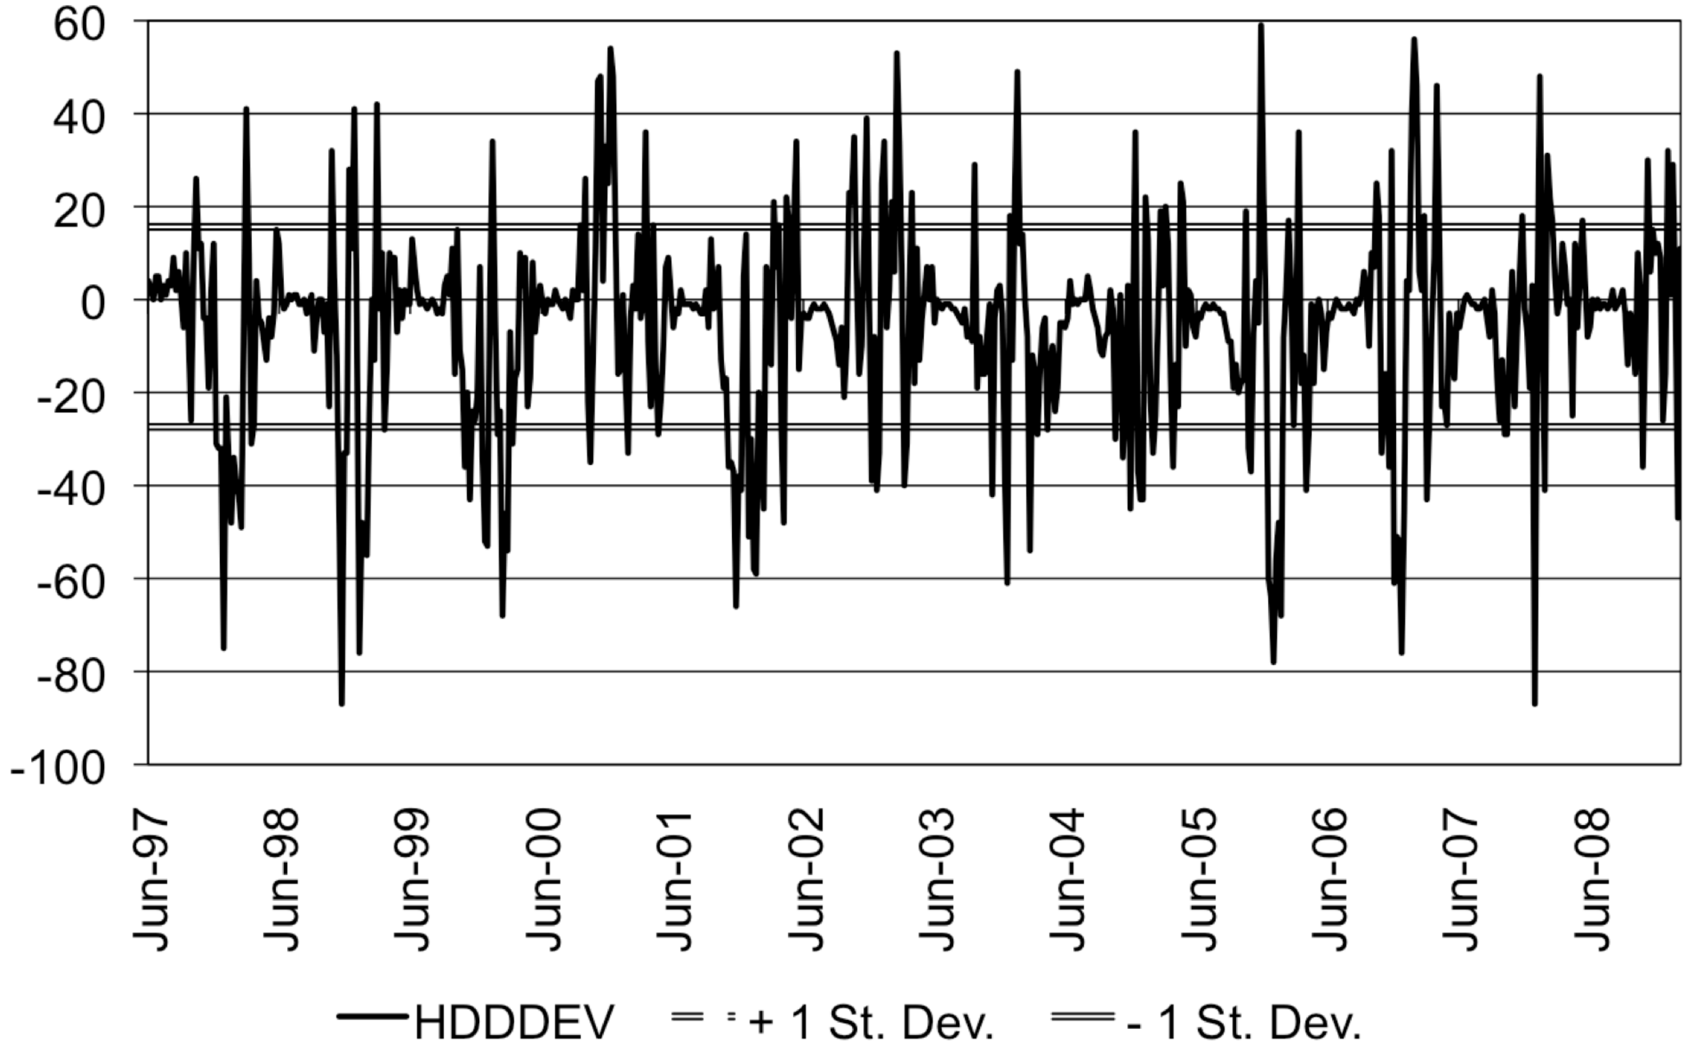

Figure A4. US Average Deviation from Normal Cooling Degree Days (CDDDEV)

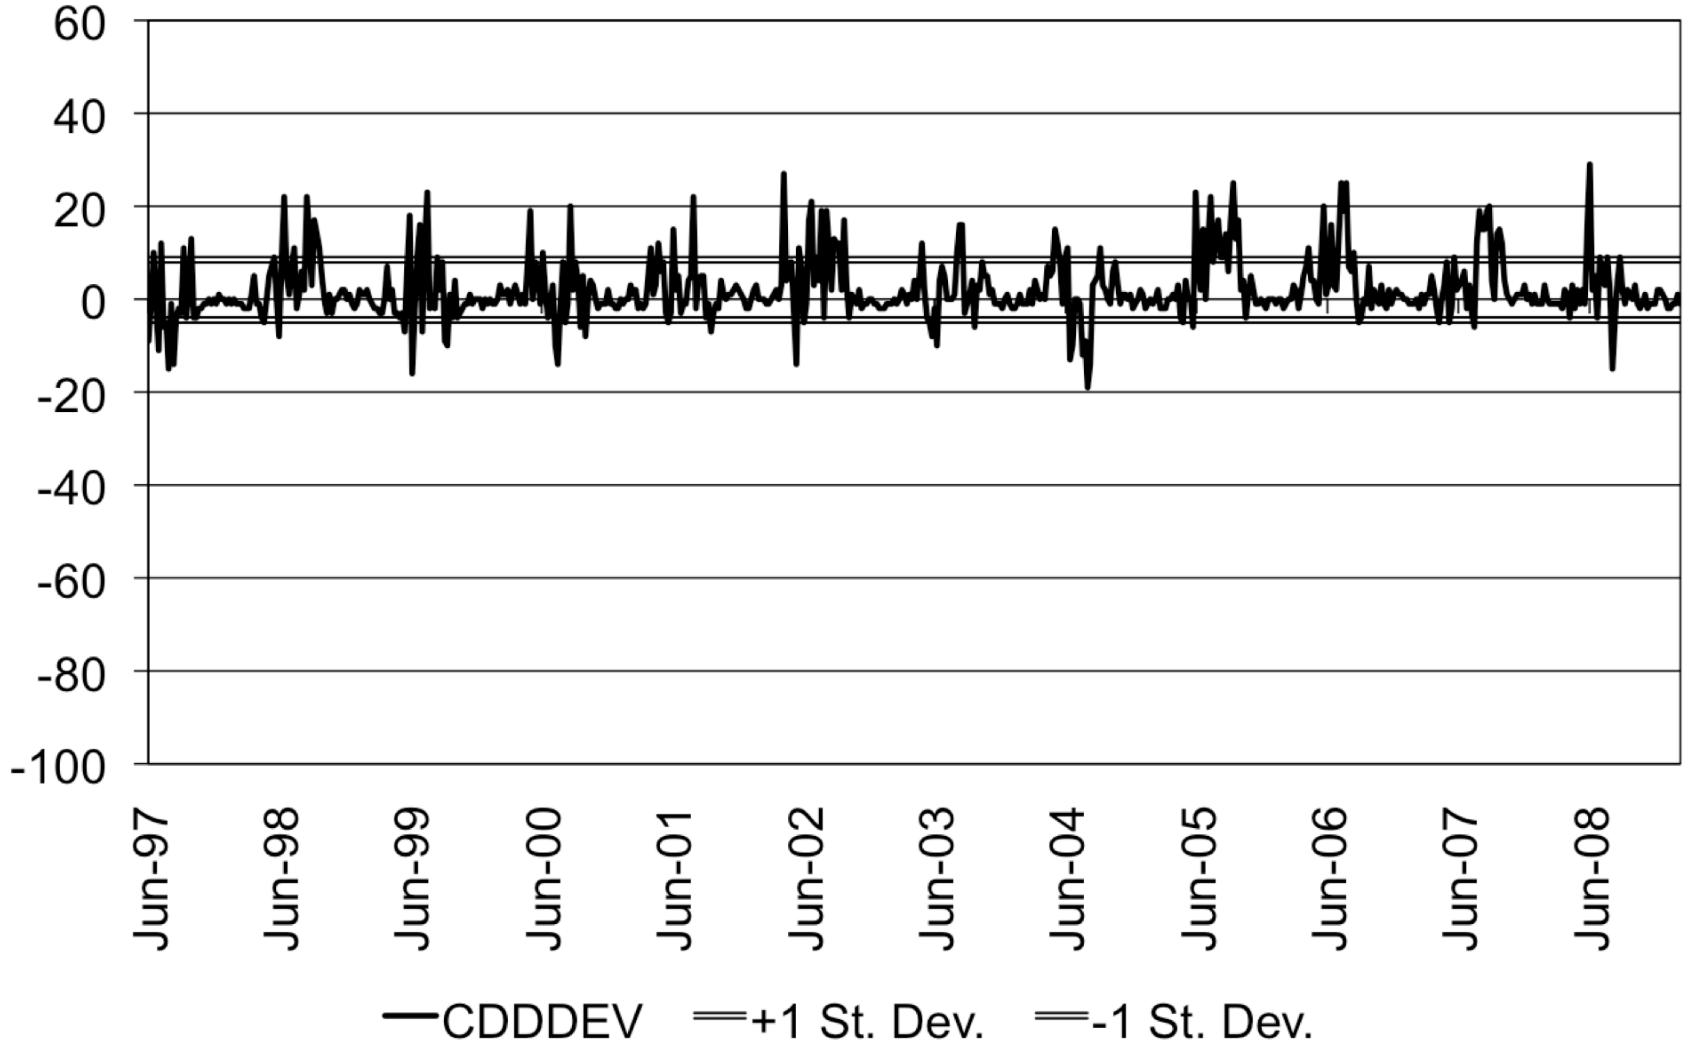

Figure A5. Shut-in Natural Gas Production Capacity in the Gulf of Mexico (SHUTIN), mmcf

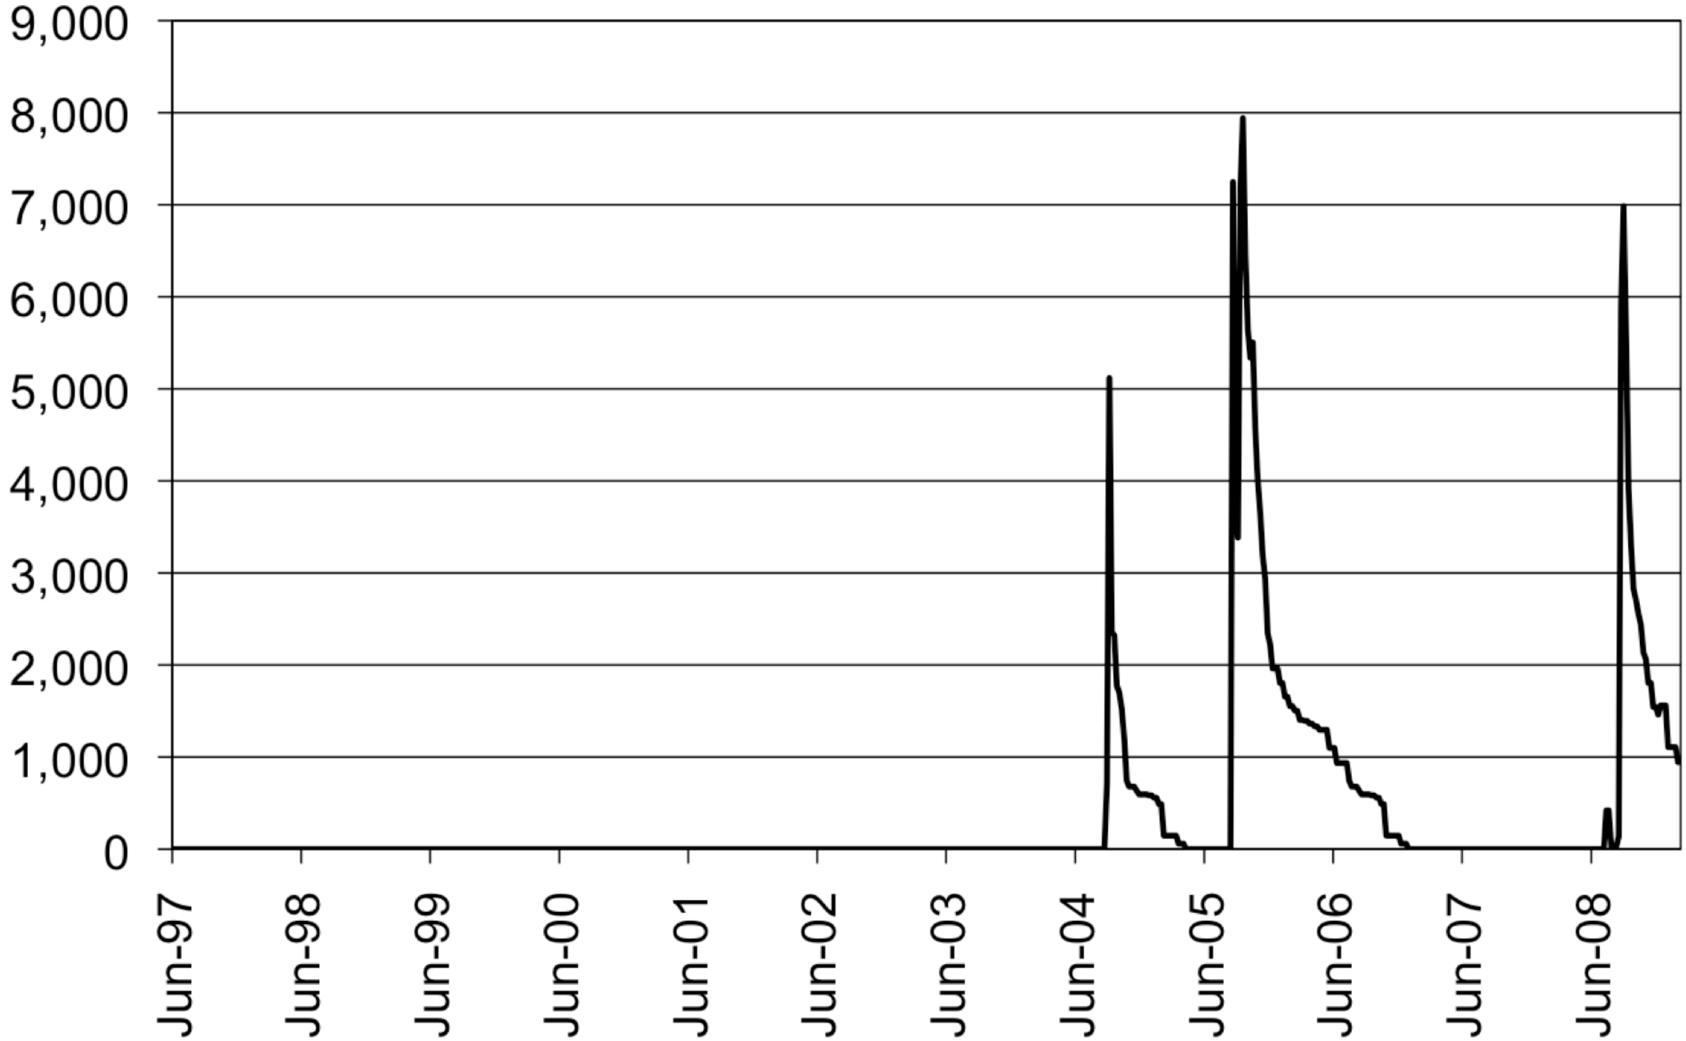

Figure A6. Differential from Running 5-year Average US Natural Gas Storage Levels (STORDIFF), Bcf

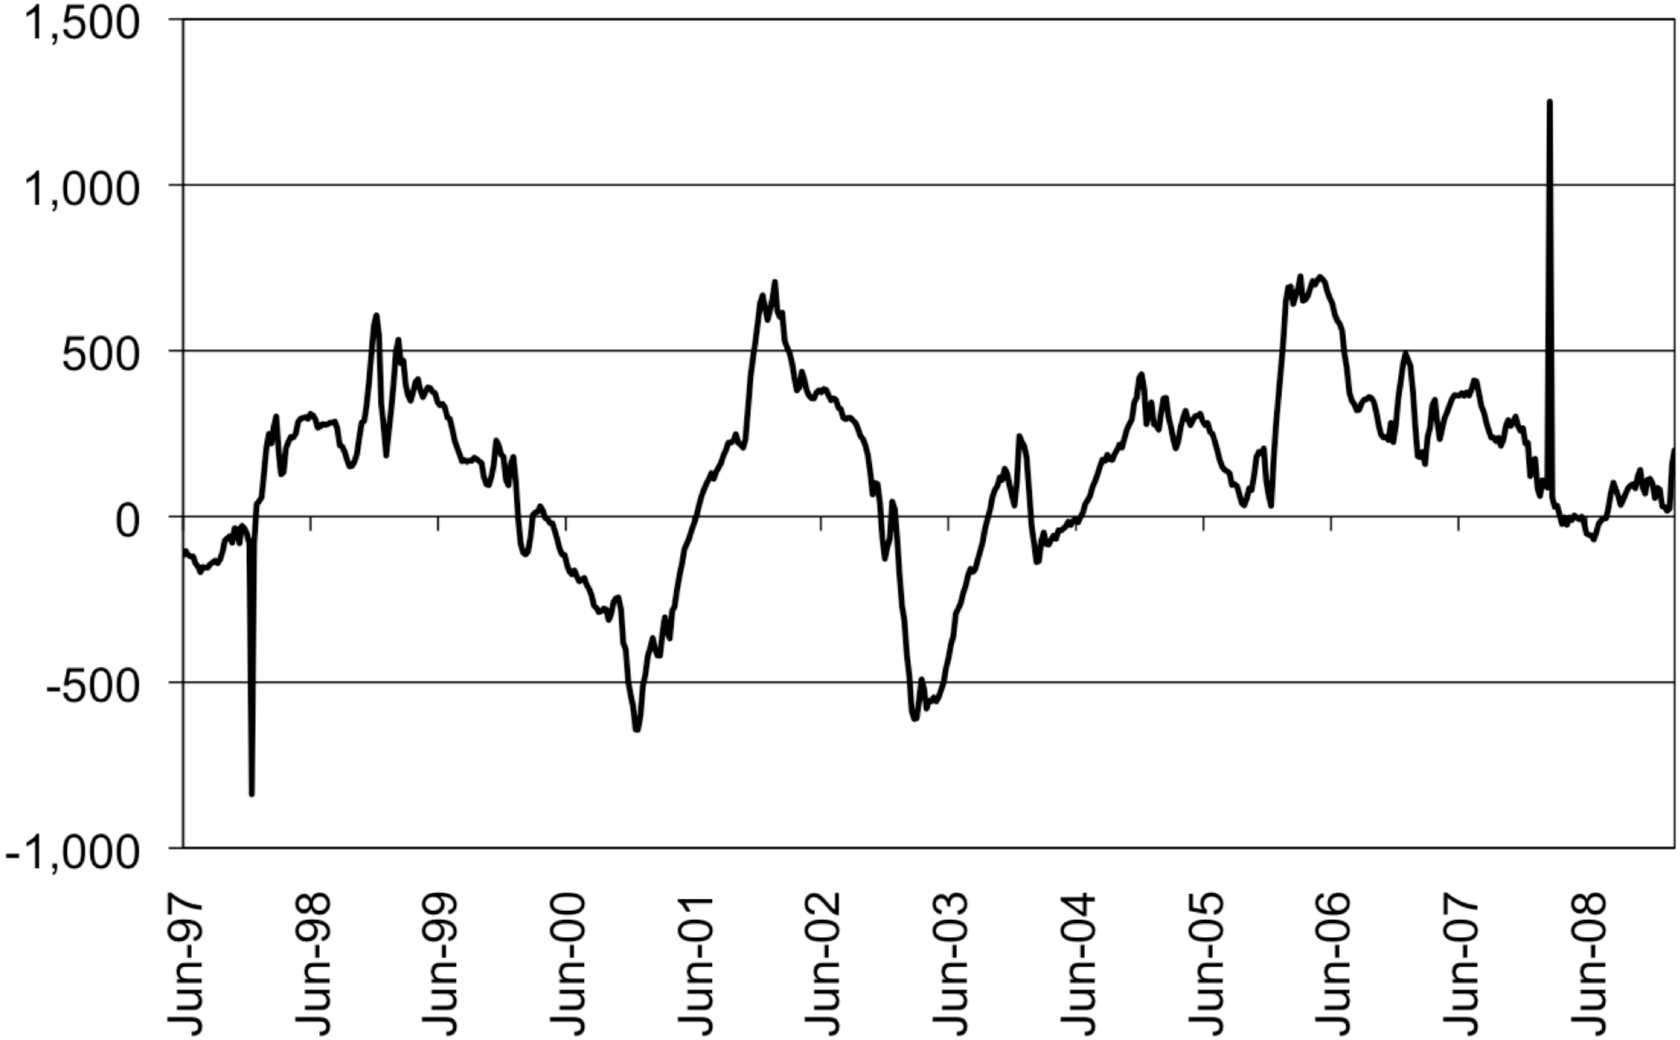

Table A.1. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 6/13/97 - 2/20/09

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.441	13.1%	-24.984	0%
lnwti	-1.477	54.5%	-27.185	0%
hdd	-3.786	0.3%	NA	
hdddev	-15.358	0.0%	NA	
cdd	-3.611	0.6%	NA	
cdddev	-16.675	0.0%	NA	
stordiff	-2.716	7.1%	NA	
shutin	-6.148	0.0%	NA	

Table A.1.2. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 6/13/97 - 2/20/09

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.441	13.1%	-24.984	0%
lnwti	-1.477	54.5%	-27.185	0%

Table A2. Phillips-Perron Tests

Dates of analysis (weekly): 6/13/97 - 2/20/09

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-9.103	-2.262	18.4%	-558.254	-25.164	0%
lnwti	-3.204	-1.399	58.3%	-664.208	-27.207	0%
hdd	-33.882	-4.111	0%	NA	NA	
hdddev	-331.427	-15.223	0%	NA	NA	
cdd	-40.339	-4.499	0%	NA	NA	
cdddev	-430.649	-17.27	0%	NA	NA	
stordiff	-15.161	-2.802	5.8%	NA	NA	
shutin	-64.148	-5.842	0%	NA	NA	

Table A2.2. Phillips-Perron Tests						
Dates of analysis (weekly): 6/13/97 - 2/20/09						
Variable	**Levels Z(rho)**	**Levels Z(t)**	**Significance Z(t) %**	**1st Differences Z(rho)**	**1st Differences Z(t)**	**Significance Z(t) %**
lnhh	-9.103	-2.262	18.4%	-558.254	-25.164	0%
lnwti	-3.204	-1.399	58.3%	-664.208	-27.207	0%

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

Dates of analysis (weekly): 6/13/97 - 2/20/09

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwatrz Bayesian Information Criterion (SBIC)
0	-379.52			0.013	1.31611	1.35616	1.41897
1	1435.07	3629.200	0.000	3.00E-05	-4.73935	-4.68786	-4.60710 *
2	1442.72	15.317	0.004	3.00E-05	-4.75159	-4.68866 *	-4.58995
3	1448.04	10.633	0.031	2.90E-05	-4.75599	-4.68162	-4.56497
4	1452.90	9.722	0.045	2.90E-05	-4.75887	-4.67306	-4.53846
5	1453.96	2.111	0.715	3.00E-05	-4.74902	-4.65177	-4.49922
6	1457.71	7.505	0.111	3.00E-05	-4.74820	-4.63950	-4.46900
7	1460.33	5.232	0.264	3.00E-05	-4.74357	-4.62342	-4.43499
8	1466.19	11.732	0.019	3.00E-05	-4.74981	-4.61822	-4.41184
9	1470.53	8.683	0.070	3.00E-05	-4.75095	-4.60792	-4.38359
10	1477.22	13.364 *	0.010	2.90E-05 *	-4.75992 *	-4.60545	-4.36317
11	1480.10	5.771	0.217	2.90E-05	-4.75619	-4.59028	-4.33006
12	1481.42	2.647	0.619	3.00E-05	-4.74724	-4.56989	-4.29172

Table A4. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 2/20/09

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	1467.331		31.621	29.811	-4.358	-4.582	-4.724
1	1482.236	0.048	1.809**	1.809	-4.376***	-4.613***	-4.764
2	1483.141	0.003			-4.368	-4.610	-4.764

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

Table A5. VECM Model including Exogenous Variables (6/13/97-2/20/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	26	0.0916	0.1411	94.2729	0.0000
D-lnwti	26	0.0576	0.1143	74.0499	0.0000

Long-Term Variables:	Values	P-Values	
β	0.7342	0.0000	**
ɣ	-1.0493		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0031	0.832	
cointegrating term (t-1) (α)	-0.0828	0.000	**
ΔPHH(t−1)	-0.0559	0.195	
ΔPHH(t−2)	-0.0710	0.096	+
ΔPHH(t−3)	-0.0931	0.026	*
ΔPHH(t−4)	-0.0590	0.156	
ΔPHH(t−5)	-0.0740	0.074	+
ΔPHH(t−6)	0.0267	0.518	
ΔPHH(t−7)	-0.0631	0.124	
ΔPHH(t−8)	0.0687	0.093	+
ΔPHH(t−9)	-0.1012	0.014	*
ΔPWTI(t−1)	0.0376	0.574	
ΔPWTI(t−2)	0.0370	0.583	
ΔPWTI(t−3)	0.0562	0.407	
ΔPWTI(t−4)	-0.0185	0.783	
ΔPWTI(t−5)	-0.0101	0.882	
ΔPWTI(t−6)	0.0876	0.197	
ΔPWTI(t−7)	0.0983	0.149	
ΔPWTI(t−8)	0.0519	0.460	
ΔPWTI(t−9)	0.1077	0.121	
HDD(t)	1.14E-04	0.191	
HDDDEV(t)	1.02E-03	0.000	**
CDD(t)	-4.23E-04	0.116	
CDDDEV(t)	3.43E-03	0.000	**
STORAGE DIFF(t)	-2.97E-05	0.076	+
SHUT IN(t)	4.49E-06	0.263	

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	24.13	0.0041	**
Lagged WTI	7.04	0.6333	
Lagged HH & WTI	30.28	0.0348	*
Exogenous Vars.	51.15	0.0000	**
Exog + HH Lag	71.71	0.0000	**
Exog + WTI Lag	54.61	0.0000	**
Lagged + Exogs	75.86	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0081	0.380	
cointegrating term (t-1) (α)	0.0318	0.011	*
ΔPHH(t−1)	0.0733	0.007	**
ΔPHH(t−2)	0.0164	0.540	
ΔPHH(t−3)	-0.0191	0.466	
ΔPHH(t−4)	-0.0001	0.997	
ΔPHH(t−5)	0.0101	0.698	
ΔPHH(t−6)	0.0308	0.235	
ΔPHH(t−7)	0.0163	0.526	
ΔPHH(t−8)	0.0006	0.982	
ΔPHH(t−9)	-0.0283	0.273	
ΔPWTI(t−1)	-0.1314	0.002	**
ΔPWTI(t−2)	-0.1024	0.016	*
ΔPWTI(t−3)	0.0891	0.036	*
ΔPWTI(t−4)	-0.0195	0.645	
ΔPWTI(t−5)	0.0357	0.401	
ΔPWTI(t−6)	-0.0625	0.143	
ΔPWTI(t−7)	-0.0818	0.056	+
ΔPWTI(t−8)	0.1035	0.019	*
ΔPWTI(t−9)	0.1126	0.010	**
HDD(t)	4.45E-05	0.419	
HDDDEV(t)	-1.66E-04	0.195	
CDD(t)	1.97E-04	0.244	
CDDDEV(t)	3.70E-04	0.418	
STORAGE DIFF(t)	2.44E-05	0.021	*
SHUT IN(t)	-7.23E-06	0.004	**

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	11.28	0.2567	
Lagged WTI	44.28	0.0000	**
Lagged HH & WTI	50.72	0.0001	**
Exogenous Vars.	17.61	0.0073	**
Exog + HH Lag	30.37	0.0107	*
Exog + WTI Lag	62.34	0.0000	**
Lagged + Exogs	70.24	0.0000	**

Table A6. Conditional ECM including Exogenous Variables (6/13/97-2/20/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	18	0.0899	0.1606	114.8139	0.0000

Long-Term Variables:	Values	P-Values	
β	0.7342	0.0000	**
γ	-1.0493		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0015	0.9160	
cointegrating term (t-1) (α)	-0.0952	0.0000	**
ΔPHH(t−1)	-0.0610	0.1350	
ΔPHH(t−2)	-0.0566	0.1550	
ΔPHH(t−3)	-0.0723	0.0640	+
ΔPHH(t−4)	-0.0536	0.1690	
ΔPHH(t−5)	-0.0747	0.0540	+
ΔPHH(t−6)	0.0334	0.3890	
ΔPHH(t−7)	-0.0419	0.2750	
ΔPHH(t−8)	0.0864	0.0240	*
ΔPHH(t−9)	-0.0816	0.0340	*
ΔPWTI	0.2872	0.0000	**
HDD(t)	7.66E-05	0.3560	
HDDDEV(t)	1.05E-03	0.0000	**
CDD(t)	-4.70E-04	0.0690	+
CDDDEV(t)	3.26E-03	0.0000	**
STORAGE DIFF(t)	-3.37E-05	0.0370	*
SHUT IN(t)	5.12E-06	0.1820	

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	24.58	0.0035	**
Exogenous Vars.	54.39	0.0000	**
Exog + WTI	72.44	0.0000	**
Lagged + Exogs	76.49	0.0000	**
Lag + Exog + WTI	95.16	0.0000	**

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

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 (6/13/97-2/20/09 weekly)	0	0	20.0	0.0
	0	1	0.0	1.8
	1	4	0.0	5.0
	2	8	0.0	10.0
	12	52	0.0	13.2
Ramberg-Parsons Conditional ECM (6/13/97-2/20/09 weekly)	0	0	20.0	5.4
	0	1	0.0	5.7
	1	4	0.0	6.5
	2	8	0.0	8.4
	12	52	0.0	13.2

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 (6/13/97-2/20/09 weekly)	0	0	20.0	0.0
	0	1	-16.7	1.8
	1	4	0.0	0.0
	2	8	0.0	0.9
	12	52	0.0	0.1
Ramberg-Parsons Conditional ECM (6/13/97-2/20/09 weekly)	0	0	20.0	5.4
	0	1	-16.7	0.6
	1	4	0.0	0.5
	2	8	0.0	1.1
	12	52	0.0	0.3

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	$0.00054	0.356	35.60%	272	0	79.83	$0.04
HDDDEV	$0.00736	0.000	0.00%	59	-87	21.52	$0.16
CDD	-$0.00329	0.069	6.90%	100	0	28.23	-$0.09
CDDDEV	$0.02288	0.000	0.00%	29	-19	6.48	$0.15
STORDIFF	-$0.00024	0.037	3.70%	1251	-838	278.61	-$0.07
SHUTIN	$0.00004	0.182	18.20%	7941	0	1036.88	$0.04
ΔPHH,t-1	-$0.41422	0.135	13.50%	0.480	-0.569	0.096	-$0.04
ΔPHH,t-2	-$0.38494	0.155	15.50%	0.480	-0.569	0.096	-$0.04
ΔPHH,t-3	-$0.48801	0.064	6.40%	0.480	-0.569	0.096	-$0.05
ΔPHH,t-4	-$0.36556	0.169	16.90%	0.480	-0.569	0.096	-$0.04
ΔPHH,t-5	-$0.50365	0.054	5.40%	0.480	-0.569	0.096	-$0.05
ΔPHH,t-6	$0.23774	0.389	38.90%	0.480	-0.569	0.096	$0.02
ΔPHH,t-7	-$0.28739	0.275	27.50%	0.480	-0.569	0.096	-$0.03
ΔPHH,t-8	$0.63168	0.024	2.40%	0.480	-0.569	0.096	$0.06
ΔPHH,t-9	-$0.54861	0.034	3.40%	0.480	-0.569	0.096	-$0.05
ΔPWTI,t	$2.32879	0.000	0.00%	0.359	-0.312	0.060	$0.12

Table A9. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 6/13/97 - 12/16/05

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-1.426	57.0%	-20.851	0%
lnwti	-0.966	76.6%	-22.333	0%
hdd	-3.076	3%	NA	NA
hdddev	-12.86	0%	NA	NA
cdd	-3.219	2%	NA	NA
cdddev	-15.174	0%	NA	NA
stordiff	-2.509	11.3%	NA	NA
shutin	-5.355	0%	NA	NA

Table A10. Phillips-Perron Tests

Dates of analysis (weekly): 6/13/97 - 12/16/05

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-3.84	-1.082	72.2%	-377.158	-21.106	0%
lnwti	-1.862	-0.748	83.4%	-435.049	-22.533	0%
hdd	-24.338	-3.436	1%	NA	NA	NA
hdddev	-233.61	-12.708	0%	NA	NA	NA
cdd	-30.025	-3.878	0%	NA	NA	NA
cdddev	-350.56	-15.774	0%	NA	NA	NA
stordiff	-11.556	-2.443	13.0%	NA	NA	NA
shutin	-49.299	-5.03	0%	NA	NA	NA

Table A11. Augmented Dickey-Fuller Tests

Dates of analysis (weekly): 12/23/05 - 2/20/09

Variable	Levels	Significance %	1st Differences	Significance %
lnhh	-2.68	7.8%	-14.017	0%
lnwti	-0.819	81.3%	-15.143	0%
hdd	-2.164	22%	NA	NA
hdddev	-8.287	0%	NA	NA
cdd	-1.643	46%	NA	NA
cdddev	-7.064	0%	NA	NA
stordiff	-1.392	58.6%	NA	NA
shutin	-3.173	2%	NA	NA

Table A12. Phillips-Perron Tests

Dates of analysis (weekly): 12/23/05 - 2/20/09

Variable	Levels Z(rho)	Levels Z(t)	Significance Z(t) %	1st Differences Z(rho)	1st Differences Z(t)	Significance Z(t) %
lnhh	-15.345	-2.89	4.7%	-198.771	-13.966	0%
lnwti	-2.14	-0.754	83.2%	-209.287	-14.989	0%
hdd	-9.11	-2.182	21%	NA	NA	NA
hdddev	-102.86	-8.422	0%	NA	NA	NA
cdd	-9.322	-2.162	22%	NA	NA	NA
cdddev	-76.751	-7.041	0%	NA	NA	NA
stordiff	-6.081	-1.824	36.9%	NA	NA	NA
shutin	-19.186	-3.25	2%	NA	NA	NA

Table A13. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 6/13/97 - 12/16/05

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwatrz Bayesian Information Criterion (SBIC)
0	-44.36			0.004	0.26956	0.32151	0.40117
1	1070.94	2230.600	0.000	2.60E-05	-4.86345	-4.79665 *	-4.69423 *
2	1074.77	7.665	0.105	2.60E-05	-4.86268	-4.78103	-4.65585
3	1084.36	19.172	0.001	2.60E-05	-4.88848	-4.79199	-4.64405
4	1086.84	4.960	0.291	2.60E-05	-4.88146	-4.77012	-4.59942
5	1091.90	10.124	0.038	2.60E-05	-4.88636	-4.76018	-4.56672
6	1095.94	8.084	0.089	2.60E-05	-4.88656	-4.74553	-4.52931
7	1101.30	10.727 *	0.030	2.60E-05 *	-4.89286 *	-4.73698	-4.49800
8	1102.97	3.335	0.503	2.60E-05	-4.88208	-4.71136	-4.44962
9	1103.37	0.793	0.939	2.60E-05	-4.86544	-4.67988	-4.39537
10	1107.43	8.124	0.087	2.60E-05	-4.86572	-4.66532	-4.35806
11	1109.40	3.940	0.414	2.70E-05	-4.85635	-4.64110	-4.31108
12	1112.00	5.195	0.268	2.70E-05	-4.84987	-4.61977	-4.26699

Table A14. Selection Order Criteria Tests (fitting a Vector Autoregression)

Dates of analysis (weekly): 12/23/05 - 2/20/09

No. of Lags	Log Likelihood	Likelihood Ratio	p-value	Final Prediction Error (FPE)	Akaike's Information Criterion (AIC)	Hannan-Quinn Information Criterion (HQIC)	Schwatrz Bayesian Information Criterion (SBIC)
0	71.06			0.002	-0.69162	-0.58464	-0.42808
1	394.01	645.910	0.000	3.60E-05	-4.55775	-4.42021	-4.21892 *
2	400.43	12.832	0.012	3.50E-05	-4.58703	-4.41892	-4.17291
3	409.27	17.685	0.001	3.30E-05	-4.64573	-4.44705 *	-4.15631
4	413.93	9.310	0.054	3.30E-05 *	-4.65367 *	-4.42443	-4.08895
5	415.46	3.060	0.548	3.40E-05	-4.62373	-4.36392	-3.98372
6	418.56	6.205	0.184	3.40E-05	-4.61285	-4.32248	-3.89754
7	419.46	1.799	0.773	3.50E-05	-4.57527	-4.25433	-3.78466
8	423.78	8.641	0.071	3.50E-05	-4.57915	-4.22765	-3.71325
9	427.04	6.524	0.163	3.60E-05	-4.57020	-4.18814	-3.62901
10	432.48	10.878 *	0.028	3.50E-05	-4.58765	-4.17502	-3.57116
11	433.99	3.019	0.555	3.60E-05	-4.55746	-4.11427	-3.46567
12	434.50	1.021	0.907	3.80E-05	-4.51516	-4.04140	-3.34808

Table A15. Johansen Tests for Cointegration

Logged Prices - including exogenous variables, 6/13/97 to 12/16/05

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	1100.967		32.049	32.013	-4.500	-4.714	-4.854
1	1116.973	0.070	0.036**	0.036	-4.531***	-4.762***	-4.913
2	1116.991	0.000			-4.517	-4.754	-4.909

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

Table A16. Johansen Tests for Cointegration

Logged Prices - including exogenous variables,12/23/05 to 2/20/09

Max. Rank (h_0=p)	Log Likelihood	Eigenvalue	Trace Statistic	Max Statistic	SBIC	HQIC	AIC
0	402.636		22.583	22.578	-4.076	-4.367	-4.565
1	413.925	0.128	0.006**	0.006	-4.120***	-4.444***	-4.666
2	413.927	0.000			-4.089	-4.424	-4.654

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

Table A17. VECM Model including Exogenous Variables (6/13/97-12/16/05)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	20	0.0912	0.1821	93.05402	0.0000
D-lnwti	20	0.0538	0.1128	53.1217	0.0001

Long-Term Variables:	Values	P-Values	
β	0.9860	0.0000	**
ɣ	-2.0142		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0015	0.925	
cointegrating term (t-1) (α)	-0.1375	0.000	**
ΔPHH(t−1)	-0.0242	0.623	
ΔPHH(t−2)	-0.1456	0.003	**
ΔPHH(t−3)	-0.0211	0.659	
ΔPHH(t−4)	-0.0911	0.058	+
ΔPHH(t−5)	0.0141	0.769	
ΔPHH(t−6)	-0.0089	0.852	
ΔPWTI(t−1)	-0.0975	0.249	
ΔPWTI(t−2)	-0.0041	0.962	
ΔPWTI(t−3)	-0.0703	0.406	
ΔPWTI(t−4)	-0.0610	0.473	
ΔPWTI(t−5)	-0.1299	0.122	
ΔPWTI(t−6)	0.0703	0.397	
HDD(t)	2.49E-04	0.014	*
HDDDEV(t)	1.41E-03	0.000	**
CDD(t)	1.22E-05	0.968	
CDDDEV(t)	3.03E-03	0.000	**
STORAGE DIFF(t)	-2.83E-05	0.146	
SHUT IN(t)	1.15E-05	0.014	*

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	11.44	0.0757	+
Lagged WTI	4.86	0.5622	
Lagged HH & WTI	19.43	0.0787	+
Exogenous Vars.	58.91	0.0000	**
Exog + HH Lag	70.12	0.0000	**
Exog + WTI Lag	61.71	0.0000	**
Lagged + Exogs	74.66	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	-0.0158	0.102	
cointegrating term (t-1) (α)	0.0135	0.385	
ΔPHH(t−1)	0.0576	0.048	*
ΔPHH(t−2)	0.0207	0.473	
ΔPHH(t−3)	-0.0363	0.199	
ΔPHH(t−4)	0.0295	0.299	
ΔPHH(t−5)	-0.0117	0.679	
ΔPHH(t−6)	0.0444	0.114	
ΔPWTI(t−1)	-0.0530	0.289	
ΔPWTI(t−2)	-0.1824	0.000	**
ΔPWTI(t−3)	0.0905	0.070	+
ΔPWTI(t−4)	-0.1201	0.017	*
ΔPWTI(t−5)	0.0720	0.147	
ΔPWTI(t−6)	-0.1329	0.007	**
HDD(t)	8.08E-05	0.179	
HDDDEV(t)	-6.88E-05	0.634	
CDD(t)	3.35E-04	0.065	+
CDDDEV(t)	4.57E-04	0.338	
STORAGE DIFF(t)	1.69E-05	0.143	
SHUT IN(t)	-2.22E-06	0.423	

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	10.3	0.1125	
Lagged WTI	29.62	0.0000	**
Lagged HH & WTI	37.47	0.0002	**
Exogenous Vars.	10.56	0.1029	
Lagged + Exogs	50.43	0.0001	*

Table A18. VECM Model including Exogenous Variables (12/23/05-2/20/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
D_lnhh	14	0.0797	0.2986	64.27627	0.0000
D-lnwti	14	0.0662	0.1586	28.4681	0.0123

Long-Term Variables:	Values	P-Values	
β	0.6782	0.0000	**
ɣ	-0.8598		

Henry Hub Effects (D_lnhh):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0078	0.777	
cointegrating term (t-1) (α)	-0.1495	0.000	**
ΔPHH(t−1)	-0.0323	0.679	
ΔPHH(t−2)	0.2540	0.001	**
ΔPHH(t−3)	-0.1963	0.009	**
ΔPWTI(t−1)	0.0106	0.912	
ΔPWTI(t−2)	-0.0599	0.541	
ΔPWTI(t−3)	0.0471	0.627	
HDD(t)	-1.75E-05	0.905	
HDDDEV(t)	1.64E-04	0.630	
CDD(t)	-1.06E-03	0.024	*
CDDDEV(t)	4.34E-03	0.002	**
STORAGE DIFF(t)	-1.03E-05	0.734	
SHUT IN(t)	-9.72E-06	0.116	

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	25.45	0.0000	**
Lagged WTI	0.86	0.8353	
Lagged HH & WTI	26.06	0.0002	**
Exogenous Vars.	16.24	0.0125	*
Exog + HH Lag	42.71	0.0000	**
Exog + WTI Lag	17.56	0.0406	*
Lagged + Exogs	43.93	0.0000	**

WTI Effects (D_lnwti):

Short-Term Variables:	Values	P-Values	
constant (a)	0.0164	0.476	
cointegrating term (t-1) (α)	0.0714	0.030	*
ΔPHH(t−1)	0.0602	0.353	
ΔPHH(t−2)	0.0547	0.377	
ΔPHH(t−3)	0.0165	0.792	
ΔPWTI(t−1)	-0.2486	0.002	**
ΔPWTI(t−2)	-0.0908	0.265	
ΔPWTI(t−3)	0.0927	0.249	
HDD(t)	-9.96E-05	0.413	
HDDDEV(t)	-1.89E-04	0.505	
CDD(t)	-1.35E-04	0.729	
CDDDEV(t)	1.99E-04	0.861	
STORAGE DIFF(t)	3.12E-05	0.215	
SHUT IN(t)	-1.44E-05	0.005	**

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	1.41	0.7035	
Lagged WTI	12.43	0.0060	**
Lagged HH & WTI	12.91	0.0445	*
Exogenous Vars.	10.97	0.0892	+
Lagged + Exogs	22.49	0.0324	

Table A19. Conditional ECM Including Exogenous Variables (6/13/97 - 12/16/05)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	15	0.0893	0.2058	113.504	0.0000

Long-Term Variables:	Values	P-Values	
β	0.9860	0.0000	**
γ	-2.0142		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0051	0.7480	
cointegrating term (t-1) (α)	-0.1361	0.0000	**
ΔPHH(t-1)	-0.0550	0.2340	
ΔPHH(t-2)	-0.1515	0.0010	**
ΔPHH(t-3)	-0.0216	0.6260	
ΔPHH(t-4)	-0.1082	0.0150	*
ΔPHH(t-5)	-0.0068	0.8770	
ΔPHH(t-6)	-0.0199	0.6520	
ΔPWTI	0.3367	0.0000	**
HDD(t)	2.09E-04	0.0320	*
HDDDEV(t)	1.40E-03	0.0000	**
CDD(t)	-1.42E-04	0.6290	
CDDDEV(t)	2.82E-03	0.0000	**
STORAGE DIFF(t)	-3.37E-05	0.0720	+
SHUT IN(t)	1.22E-05	0.0070	**

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	15.85	0.0146	*
Exogenous Vars.	62.80	0.0000	**
Exog + WTI	80.58	0.0000	**
Lagged + Exogs	76.56	0.0000	**
Lag + Exog + WTI	94.57	0.0000	**

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

Table A20. Conditional ECM Including Exogenous Variables (12/23/05-2/20/09)

Equation	Parms	RMSE	R-sq	chi2	P>chi2
dlnhh	12	0.0784	0.3081	73.4754	0.0000

Long-Term Variables:	Values	P-Values	
β	0.6782	0.0000	**
γ	-0.8598		

Henry Hub Effects (dlnhh)

Short-Term Variables:	Values	P-Values	
constant (a)	0.0046	0.8610	
cointegrating term (t-1) (α)	-0.1642	0.0000	**
ΔPHH(t-1)	-0.0413	0.5730	
ΔPHH(t-2)	0.2458	0.0000	**
ΔPHH(t-3)	-0.1959	0.0050	**
ΔPWTI	0.1851	0.0380	*
HDD(t)	8.27E-07	0.9950	
HDDDEV(t)	1.78E-04	0.5770	
CDD(t)	-1.04E-03	0.0200	*
CDDDEV(t)	4.35E-03	0.0010	**
STORAGE DIFF(t)	-1.53E-05	0.5900	
SHUT IN(t)	-7.53E-06	0.1840	

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

Joint Significance:

Variable	Chi2 Stat	P-Value	
Lagged HH	27.78	0.0000	**
Exogenous Vars.	17.80	0.0067	**
Exog + WTI	22.92	0.0018	**
Lagged + Exogs	46.57	0.0000	**
Lag + Exog + WTI	52.32	0.0000	**