Revisiting the Drivers of Natural Gas Prices.

A replication study of Brown & Yücel (*The Energy Journal*, 2008)

Gavin Roberts*

International Journal for Re-Views in Empirical Economics, Volume 3, 2019-2, DOI: 10.18718/81781.11

JEL: D72, F18, F59, Q56
Keywords: energy, crude oil, natural gas, cointegration, replication study

Data Availability: The variables used in this replication are analogous to the variables used in the original Brown and Yücel (2008) paper. A detailed description is given in Section 3.1 of this paper. The data set and Stata do file are available from the website of the journal www.iree.eu.

Please Cite As: Roberts, Gavin (2019). Revisiting the Drivers of Natural Gas Prices. A replication study of Brown & Yücel (The Energy Journal, 2008). *International Journal for Re-Views in Empirical Economics*, Vol 3(2019-2). DOI: 10.18718/81781.11

Abstract

This paper replicates the analysis in the paper “What Drives Natural Gas Prices?" by Stephen P.A. Brown and Mine K. Yücel. The replication confirms the results of that analysis: a long-run relationship existed between natural-gas prices and crude-oil prices during the period from June 1997 to June 2007. This relationship was primarily driven by crude-oil prices, as natural-gas prices adjusted to deviations from the long-run relationship. Controlling for exogenous covariates related to weather, seasonality, and supply disruptions strengthen the price relationship between these two commodities. When the sample is expanded to include data generated as recently as June 2017, evidence of the long-run relationship disappears completely. I posit that this results from increased U.S. natural-supply associated with the “shale revolution”.

*Weber State University, Email: gavinroberts@weber.edu, Phone: (720) 317-3086. I have no competing interests associated with this article to declare.
Acknowledgements: The author thanks two anonymous referees and the editor for wonderful comments that improved the paper, and thanks Benjamin Gilbert for his exceptional instruction on the methods applied in this paper at the University of Wyoming in 2013.

Received July 16, 2018; Revised January 1, 2019; Accepted February 8, 2019; Published March 12, 2019.

1 Introduction

In January of 2008 the paper “What Drives Natural Gas Prices?" by Stephen P.A. Brown and Mine K. Yücel was published in the *Energy Journal* (Brown and Yücel, 2008). This paper analyzes the drivers of natural-gas prices by estimating a long-run relationship between gas prices and oil prices, and buttressing the model by controlling for natural-gas-market specific factors. The paper has received 294 citations as of this writing according to Google Scholar.[1] The majority of these citations have occurred since 2013, and include eight citations in 2019 and 30 citations in 2018, indicating continued strong interest in the analysis of the impact of crude-oil prices and other factors on natural-gas prices. In this paper, I replicate the analysis of Brown and Yücel (2008) by re-collecting the data, and expanding the sample of the analysis to include data generated as recently as June 2017.

The articles that cite Brown and Yücel (2008) relate to a diverse set of research questions including reexaminations of the relationship between gas and oil prices, the policy implications of links between gas and oil markets, and other topics. Some articles point out that the long-run relationship between natural-gas prices and crude-oil prices appears to be rather weak, or to have disappeared altogether, especially when data from post-2008 are included in the analysis, as increasing supplies of natural gas associated with the shale revolution began to impact the market (Erdős, 2012; Ramberg and Parsons, 2012). Brigida (2014) rediscovered a strong relationship between gas prices and oil prices by allowing shifts in the cointegrating vector, but Caporin and Fontini (2017) find mixed evidence of such a long run relationship in a model that controls for new production of natural gas from shale reservoirs. One recent article found that shale gas has significantly depressed crude oil price in contrast to the more popular view that crude oil prices drive natural gas prices (Liu and Li, 2018). Others have discovered similar links between crude-oil and natural-gas prices in different markets such as the German natural-gas market (Nick and Thoenes, 2014). Importantly, many of the articles that cite Brown and Yücel (2008) discuss policy implications related to the relationship between crude-oil and natural-gas prices, so changes to this relationship must be examined in order to update such policy implications, for example see Pettersson et al. (2012) and Brown et al (2010).

There are several benefits associated with the replication of the analysis of Brown and Yücel (2008). First, it offers an independent verification of the results over the same sample period. Indeed, the results of my replication are qualitatevely similar to the results presented in Brown and Yücel (2008). Small quantitative differences likely result from differences in the data collection process, which are discussed in more detail below. Second, some of the data from the original analysis were collected from a proprietary dataset, while this replication uses data that was collected from the original public data sources. Therefore, the data can be disseminated to interested scholars for further replication and analyses. Third, a cataclysmic shift in natural gas markets, and energy markets more generally, was beginning to occur just at the end of the sample period of the Brown and Yücel (2008) analysis, as the shale revolution was beginning to take hold and to drastically increase the supply of natural gas in the United States. This replication indicates that the relationship between crude-oil prices and natural-gas prices espoused in Brown and Yücel (2008) was knocked out of whack after the end of the original sampling period, and I suspect that the shale revolution was the primary culprit in this relational change.

[1] See `scholar.google.com/citations?view_op=view_citation&citation_for_view=XFI42IoAAAAJ:qjMakFHDy7sC`.

The format of this replication follows the format of Brown and Yücel (2008) closely for ease of comparison with the original paper. In Section 2, industry rules-of-thumb for the relationship between gas and oil prices are discussed and analyzed in Subsection 2.1, and burner-tip parity rules are discussed and analyzed in Subsection 2.2. In Section 3, a vector-error correction model is utilized in a more advanced appraisal of the effect of crude-oil prices on natural-gas prices. Subsection 3.1 provides a detailed description of the data and the data collection process from publicly available sources for this replication study, while Subsection 3.2 presents the econometric model and associated results. Section 4 offers concluding remarks.

2 The Relationship Between Oil and Natural Gas Prices

2.1 Simple Rules of Thumb

Brown and Yücel (2008) discuss two commonly employed industry rules-of-thumb in their analysis of the relationship between crude-oil and natural-gas prices. Both rules are constant ratios of crude-oil price to natural-gas price. These are the 10-to-1 rule and the 6-to-1 rule. The 10-to-1 rule was commonly applied by industry stakeholders because it provided a relatively good fit to historical price data. The 6-to-1 rule was applied because it reflects the higher energy content of a barrel of crude oil. Natural gas is priced in U.S. dollars per million British thermal units ($/MMBtu), while crude oil is priced in dollars per barrel. There are approximately 6 MMBtu of energy content in a barrel of crude oil. The 6-to-1 rule reflects this energy-content difference.

Figure 1 is analogous to Figure 1 in Brown and Yücel (2008). It charts the actual Henry Hub natural-gas price in levels, and the natural-gas prices implied by the 10-to-1 and 6-to-1 rules given the West Texas Intermediate crude-oil price.[2] The vertical line just to the right of 2007 in Figure 1 indicates the last week of the Brown and Yücel (2008) dataset, which occurred in June 2007. As indicated by Brown and Yücel (2008), the 10-to-1 rule offers a better fit than the 6-to-1 rule until approximately 2005, when the actual Henry Hub price begins to converge to the 6-to-1 rule. Using the expanded dataset we see that the 10-to-1 rule has consistently overestimated the Henry Hub natural-gas price since 2005, while the 6-to-1 rule offered a relatively good fit until early 2009. Since early 2009, however, both the 10-to-1 and 6-to-1 rule have greatly overestimated the Henry Hub natural-gas price, even after the precipitous decrease in crude-oil prices that began in 2014.

2.2 Burner-Tip Parity Rules

The 6-to-1 rule implicitly assumes a demand-side driven long-run relationship between crude-oil prices and natural-gas prices, because it implies prices per unit of energy should be held in check by demand-side substitution – if one of the commodities is more expensive per unit of energy then consumers should shift consumption toward the other commodity keeping the difference in prices between the two commodities in check.

[2] These series are the West Texas Intermediate crude-oil price divided by 10 and divided by 6, respecitively.

Figure 1: Actual and Implied Natural Gas Prices (Levels)

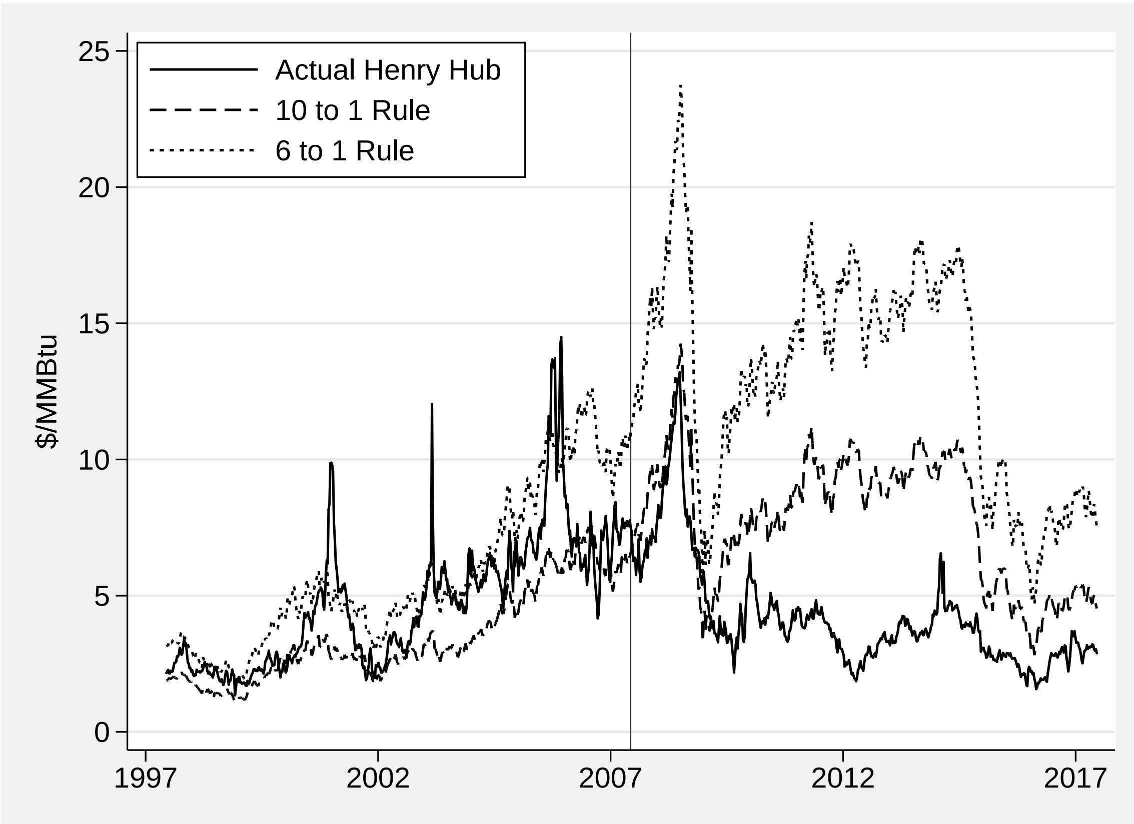

Burner-tip parity rules represent a more nuanced view of demand-side competition between crude-oil and natural-gas prices. Rather than simply positing that the ratio between crude-oil prices and natural-gas prices should be constant, burner-tip parity rules account for the fact that barrels of crude oil do not directly compete with MMBtus of natural gas at the burner tip (point of conversion to energy). Natural gas generally competes more directly with either Residual Fuel Oil (RFO) or Distillate Fuel Oil (DFO) at the burner tip, as these are the most commonly used petroleum products in the generation of electricity and/or space heating. Burner-tip parity rules also account for cost differentials between transporting natural gas to the burner tip versus transporting these petroleum products to the burner tip, where the former is generally more costly as a result of the need for pipeline capacity. Let $P_{HH,t}$ be the price of natural gas per MMBtu at the Henry Hub in Louisiana at time t (the U.S. benchmark natural-gas price), let $P_{WTI,t}$ be the price per barrel of crude oil at Cushing, Oklahoma (the U.S. benchmark crude-oil price), let $P_{BT,t}$ be the price per barrel of the petroleum product that directly competes with natural gas at the burner tip, let E_{BT} be the energy content in MMBtu of the petroleum product that directly competes with natural gas at the burner tip, and let T be the transportation cost differential (the cost of transporting the petroleum product minus the cost of transporting natural gas). Now, we can write the burner-tip parity rule as

$$P_{HH,t} = T + \frac{P_{BT,t}}{P_{WTI,t}} \frac{1}{E_{BT}} P_{WTI,t}. \tag{1}$$

The T term in equation (1) is negative if the cost of transporting natural gas from the Henry Hub to the burner tip is more than the cost of transporting the petroleum product to the burner tip, and the natural-gas price will be lower to reflect this higher transportation cost to market. Note that if T is zero, if the price of the petroleum product is equal to the WTI price of crude oil, and if the petroleum product has an energy content of 6 MMBtu then the burner-tip parity rule is the 6-to-1 rule. However, T is generally negative, the prices of RFO and DFO generally differ from the WTI price, and RFO and DFO do not have the same energy content as a barrel of crude oil, so the burner-tip parity rule represents a more nuanced approach, as pointed out by Brown and Yücel (2008). Importantly, burner-tip parity rules implicitly assume that several terms are not time varying including the transportation cost differential, the ratio of the petroleum product price to the WTI price, and the energy content of the petroleum product.

Brown and Yücel (2008) offer estimates of these terms for RFO and DFO. The energy content of a barrel of RFO is given as 6.287 MMBtu per barrel, while the energy content of a barrel of DFO is given as 5.825 MMBtu.[3] The price ratios of RFO price to WTI price and DFO price to WTI price are given as 0.85 and 1.2, respectively.[4] Finally, the transportation cost differentials are reported as -\$0.25 for RFO and -\$0.80 for DFO in Brown and Yücel (2008). Plugging these numbers into (1) and simplifying gives

$$P_{HH,t} = -0.25 + 0.1325 P_{WTI,t} \tag{2}$$

and

$$P_{HH,t} = -0.80 + 0.2060 P_{WTI,t} \tag{3}$$

The natural-gas prices implied by the RFO burner-tip parity rule and the DFO burner-tip parity rule are charted in Figure 2 along side the actual Henry Hub natural-gas price from June 1997 to June 2017. The vertical line just to the right of 2007 in Figure 1 marks the end of the original sample from Brown and Yücel (2008). Figure 2 is analogous to Figure 2 in Brown and Yücel (2008). The RFO burner-tip parity rule provides a relatively good fit to the actual Henry Hub natural-gas price before beginning to regularly overestimate the Henry Hub price just as the Brown and Yücel (2008) dataset ends in the summer of 2007. The DFO burner-tip rule regularly overestimates the Henry hub price for the entirety of the sample. The overestimation by both of the burner-tip parity rules becomes quite extreme after the Brown and Yücel (2008) sample ends, similar to the situation with the 10-to-1 and 6-to-1 rules.

The upshot of this discussion is that simple rules of thumb and burner-tip parity rules, especially the 6-to-1 rule and the RFO burner-tip parity rule, worked quite well between June 1997 and June 2007. However, these rules began to perform extremely poorly after 2007 when the shale revolution

[3] The EIA gives the same energy content for RFO in Appendix G of its 2016 Annual Energy Outlook, while it gives a slightly lower energy content for DFO of 5.778 MMBtu per barrel EIA (2016). We use the numbers from Brown and Yücel (2008) for ease of comparison.

[4] Using monthly price data from the EIA, I found that the average ratio for RFO to WTI from 1997 to 2017 was approximately 0.80 rather than 0.85, however it is closer to 0.85 from 1997 to 2007 – the original sample period. I use 0.85 in the calculations and chart to follow for ease of comparison. I found that the average ratio for DFO to WTI was almost exactly 1.2 from 1997 to 2017.

Figure 2: Actual and Implied Natural Gas Prices (Levels)

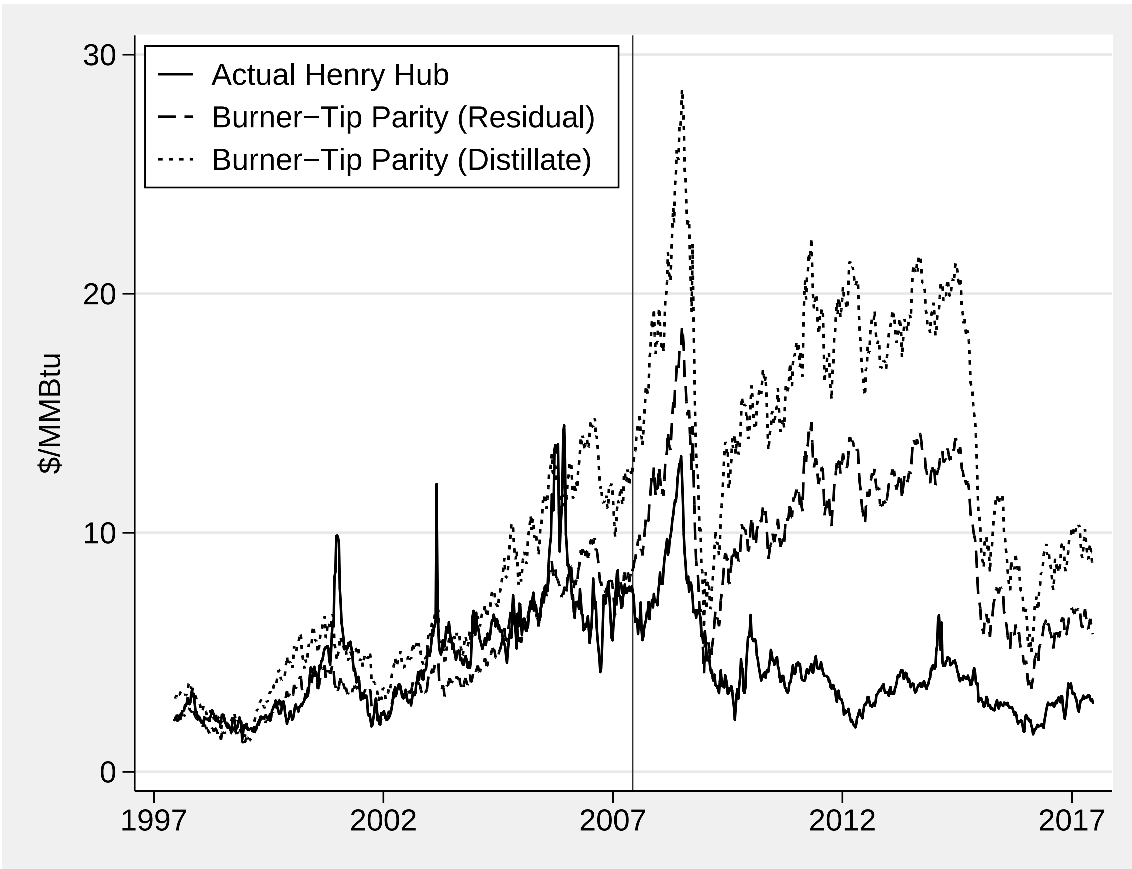

began to take hold, and the U.S. economy began to recover from the great recession. This is not to say, however, that the relationship between crude-oil prices and natural-gas prices completely broke after the Brown and Yücel (2008) sample period ended in 2007. Brown and Yücel (2008) included exogenous covariates in one of their models in an attempt to control for shocks to natural-gas prices that might be independent of crude-oil price movements. These include weather indicators and natural gas storage levels.

3 *Seasonality, Weather, Storage and Other Factors Driving Natural Gas Prices*

Opportunities for fuel switching between petroleum products and natural gas were already in decline during the Brown and Yücel (2008) sample period as electricity generation from petroleum products decreased steadily in the early-2000s, while electricity generation from natural gas continued to increase. Indeed, another paper appearing in the same issue of the *Energy Journal* studied this phenomenon and its effects on the relationship between crude-oil and natural-gas prices more explicitly by including a covariate that accounted for the capacity-weighted heat rate of petroleum-product generators and natural-gas generators (Hartley et al., 2008). According to the EIA's Electricity Data Browser[5] approximately 197,000 barrels of petroleum products (1,182,000 MMBtu) were consumed in the generation of electricity in 2001, while 82,000 barrels (492,000 MMBtu) were consumed in 2007, and 21,000 barrels (126,000 MMBtu) were consumed in 2016. On the other hand, 5,800,000 MMBtu of natural gas was burned in the generation of electricity in 2001, while 7,000,000 MMBtu were consumed in 2007, and 10,400,000 MMBtu were consumed in 2016.[6] Clearly, natural gas has played an increasingly important role in the generation of electricity in the U.S. since the turn of the century, while the role played by petroleum products has steadily declined.

Less clear at the time of the publication of Brown and Yücel (2008) was the explosive growth in natural-gas production that was about to occur in the U.S. as a result of new technologies applied to shale-gas reservoirs. The same technologies would eventually lead to explosive growth in crude-oil production, but not until much later, and not until after U.S. crude-oil production declined substantially throughout the original Brown and Yücel (2008) sample period.

Brown and Yücel (2008) were aware of the possibility that factors other than demand-side competition might lead to deviations from any long-run relationship that exists between crude-oil and natural-gas prices, and thus included a model in their analysis that attempted to control for such factors. In particular, the second model presented in their paper controls for "seasonality, extreme weather events, natural gas in storage, and disruptions to production." In particular, seasonality and extreme weather events are controlled for by using temperature data as well as deviations from normal seasonal temperatures, while natural gas in storage is used to control for a variety of possible natural-gas demand and supply shocks, and hurricane data are used to control for hurricane-related supply disruptions. A primary question in this replication is whether these control variables alone are enough to adjust for the tectonic changes that have occurred in gas and oil markets since 2007. I find that while the inclusion of these control variables does not lead to the maintenance of the long-run relationship between natural-gas and crude-oil prices, a model including these covariates continues to explain natural-gas prices quite well.

3.1 Data

The variables used in this replication are analogous to the variables used in the original Brown and Yücel (2008) paper. The primary variables of interest are weekly natural-gas prices and weekly crude-oil prices from the week ending June 13, 1997 to the week ending June 16, 2017 (the original dataset ended on the week ending June 8, 2007). Weekly spot prices of the Henry Hub natural-gas price and the WTI crude-oil price are collected from the EIA. Brown and Yücel (2008) collected these data as a weekly series from the Haver Analytics database,[7] which was originally derived from prices reported in the *Wall Street Journal*. Logged crude-oil and natural-gas prices are used

[5] See www.eia.gov/electricity/data/browser/.

[6] These calculations from barrels to MMBtu all use 6 MMBtu per barrel of petroleum product, and 1 MMBtu per thousand cubic feet (Mcf) of natural gas.

[7] See www.haver.com.

for the remainder of the analysis. Brown and Yücel (2008) logged the price series due to the large degree of volatility occurring in these series over the original sampling period.

The weather variables, used to control for the seasonal nature of natural-gas demand and natural-gas demand shocks associated with extreme temperatures, are heating degree days (HDD), deviations from normal heating degree days (HDDDEV), cooling degree days (CDD), and deviations from normal cooling degree days (CDDDEV). Heating degree days measure the magnitude of the difference between the recorded temperature in a location and 65 degrees Fahrenheit when the temperature is below 65 degrees, while cooling degree days measure the magnitude of the difference between the recorded temperature and 65 degrees Fahrenheit when the temperature is greater than 65 degrees. The idea is to measure the demand for heating or cooling, respectively, in a given location. Deviations from normal degree days measure the difference between degree days in a given week, and the average degree days for that week since 1971. The National Oceanic and Atmospheric Administration (NOAA) collects degree day data and makes it available for public use. I collected weekly observations of these degree day data directly from NOAA's website.[8] Brown and Yücel (2008) acquired their degree day data from Haver Analytics. The degree day data used in this replication is aggregated to the continental U.S., as the continental U.S. represents the primary market of interest to the original analysis. Importantly, the aggregation process applied by NOAA involves population weighting, e.g., an extremely cold day in New York City will lead to a larger addition to aggregate heating degree days than a similarly cold day in Laramie, Wyoming. The inclusion of the raw population-weighted degree day series in the model controls for seasonality in addition to controlling for weather, while deviations from normal degree days controls for the impact of weather surprises such as a polar vortex. Indeed, the impact of the large polar vortex on natural-gas prices in the U.S. during the winter of 2013 and 2014 can be seen in Figures 1 and 2.

Another control variable included in the Brown and Yücel (2008) analysis is deviation from five-year average natural-gas storage levels. Weekly working gas in storage is another series collected from the EIA, where working natural gas is natural gas that is meant to move into and out of storage facilities, as opposed to base gas, which is meant to remain permanently in storage facilities to maintain pressure and improve operability. When Brown and Yücel (2008) published their paper, the EIA provided five-year average working-gas storage levels, however, the EIA no longer provides this data. Therefore, I constructed the five-year average series myself. The first observation in the weekly working gas in storage series from EIA occurred on the week ending December 31, 1993. Therefore, the five-year average series only begins to include five years of data during the last week of 1998. For weeks between June 1997 and December 1997 only three prior observations are used in the average, while for weeks between January 1998 and December 1998 only four weeks are used. This is one departure from the original dataset that may lead to some small differences. The deviation from five-year average series is constructed by subtracting the five-year average series from the current level of working gas in storage. Above average storage levels should be expected to exert downward pressure on natural-gas prices, while below average storage levels should be expected to exert upward pressure on natural-gas prices. Natural gas storage levels reached extremely high levels relative to average in 2012, and the resulting downward impact on prices can easily be seen in Figures 1 and 2. Importantly, storage deviations from normal accumulate over time: if natural gas in storage is higher than average one week, then it can only come back to average the following week if above average demand leads to above average storage withdrawals.

[8] See www.cpc.ncep.noaa.gov/products/analysis_monitoring/cdus/degree_days/.

I will return to this point in my discussion of the results of the stationarity tests below.

The last control variable used in the analysis is production shut-in levels associated with hurricanes in the Gulf of Mexico. The Bureau of Ocean Energy Management (BOEM) distributes press releases that summarize the impact of hurricanes on crude oil and natural gas production in the Gulf of Mexico.[9] These reports include the aggregate natural gas production that was shut-in associated with Gulf hurricanes over a time period. I constructed the hurricane shut-in data series by simply dividing the aggregate shut-in production over the number of weeks affected.[10] This shut-in variable is zero for time periods in which no hurricane impact is present.

3.2 Model Specification

The first step taken by Brown and Yücel (2008), and the first step that should be taken in any cointegration analysis, is the testing of the time series of interest for the presence of unit roots. Table 1 is analogous to Table 1 in Brown and Yücel (2008), but with several additional tests. The top section of Table 1 displays the results of Augmented Dickey-Fuller Unit Root Tests (ADF tests) on the time series variables of interest. The second and third columns, labeled "Levels (BY)" and "Differences (BY)" display test statistics from the original Brown and Yücel (2008) analysis, where differences are only tested in the case that the null hypothesis of non-stationarity was not rejected in the series levels. The fourth and fifth columns, labeled "Levels (replication)" and "Differences (replication)" give the results of tests applied to my data over the Brown and Yücel (2008) sampling period, while the fifth and sixth columns labeled "Levels (full)" and "Differences (full)" give the test statistics from the updated full dataset, which ends in June 2017 rather than June 2007.

Brown and Yücel (2008) do not explicitly mention the data-generating process they are assuming under the null hypotheses in any of these ADF tests, i.e., whether the underlying autoregression contains a trend term or a drift term.[11] They also do not discuss the lag order applied in their ADF tests. Here, I use data visualization to make assumptions about the underlying data-generating process, and use Stata's *varsoc* command to determine the optimal lag order of the univariate autoregression underlying each ADF test. The results of the ADF tests on the logged price series are presented with and without a trend, while the ADF tests on the remainder of the exogenous covariates does not include a trend. The lag-order used for each ADF test is available upon request.

[9]See www.boem.gov/Press-Releases/. The Minerals Management Service of the U.S. Department of Interior was responsible for these press released when Brown and Yücel (2008) were working on their analysis.

[10]Hurricanes Gustav, Ike, Ivan, Katrina, and Rita represent the entirety of hurricanes accounted for in these data. The BOEM releases a single press release when impacts overlap, which was the case for Gustav and Ike, and for Katrina and Rita.

[11]See Hamilton (1994).

Table 1: Unit Root Tests

Variables	Levels (BY)	Differences (BY)	Levels (replication)	Differences (replication)	Levels (full)	Differences (full)
Augmented Dickey-Fuller						
logged PHH (no trend)	−1.497	−13.240**	−1.997	−15.855**	−2.921*	−22.189**
logged PHH (w/ trend)			−3.602*		−2.919	
logged PWTI (no trend)	−0.486	−7.731**	−0.703	−12.518**	−1.768	−15.255***
logged PWTI (w/trend)			−2.940		−1.488	
HDD	−8.468**		−3.466**		−6.327**	
HDDDEV	−6.689**		−12.086**		−17.244**	
CDD	−6.812**		−6.718**		−8.852**	
CDDDEV	−7.960**		−7.504**		−9.504**	
STOR DIFF	−3.406*		−2.351[1]	11.686**	−3.356*	−16.666**
SHUT IN	−4.311**		−2.845[2]		−3.925*	
Phillips-Perron						
logged PHH (no trend)			−1.899	−19.845**	−2.723	−28.038**
logged PHH (w/ trend)			−3.406		−2.716	
logged PWTI (no trend)			−0.718	−19.769**	−1.769	−27.459sym***
logged PWTI (w/trend)			−2.775		−1.461	
Zivot-Andrews						
logged PHH (break in trend & intercept)			−4.306		−5.633**	
logged PHH (break in trend only)			−3.932		−4.956**	
logged PWTI (break in trend & intercept)			−3.539		−3.775	
logged PWTI (break in trend only)			−2.661		−3.754	
Kwiatkowski-Phillips-Schmidt-Shin (KPSS)[3]						
logged PHH (no trend)			7.24 **	0.023	2.56 **	0.076
logged PHH (w/ trend)			0.202*		2.56 **	
logged PWTI (no trend)			8.14 **	0.062	10.40 **	0.172
logged PWTI (w/trend)			0.504**		2.46 **	

Note: * and ** denote significance at better than 0.05 and 0.01 levels, respectively.
[1]For this test statistic, the null hypothesis of non-stationary is rejected at the 0.1% level if a drift term is included in the ADF test, while the null is not rejected otherwise.
[2]For this test statistic, the null hypothesis of non-stationary is rejected at the 0.1% level if a drift term is included, and at the 1% level without a drift term.
[3]The null hypothesis in the KPSS tests is stationarity, so statistical significance of KPSS statistics indicates non-stationarity.

The ADF tests in Brown and Yücel (2008) failed to reject the null hypothesis of non-stationarity for both of the logged price series, and strongly reject the null hypothesis of non-stationarity in the differenced logged price series for both crude-oil and natural-gas prices. I have presented the Brown and Yücel (2008) results for the price series in the rows labeled with “no trend," as the similarity between their results and my results without a trend included in the data-generating process cause me to suspect that a trend was not used in their analysis. However, based on visualization of the crude-oil and natural-gas price series between June 1997 and June 2007, it appears that a trend should be included. When a trend is included in the Henry Hub natural-price series, the null hypothesis of non-stationarity is actually rejected at the 5% level over my subsample that matches the Brown and Yücel (2008) sample indicating the importance of choosing the correct underlying data-generating process. Small differences in the ADF test statistics are likely attributable to the use of different lag orders or small differences in the underlying data.

It becomes less clear whether a trend should be included when the ADF tests on the price series are run on the new full sample. In the full sample, the null hypothesis of non-stationarity is rejected at the 5% level for logged natural-gas prices when a trend term is excluded, while the null is not rejected when a trend is included. The ADF tests continue to indicate non-stationarity in the logged crude-oil price series whether or not a trend is included over the full sample.

The ADF tests on the exogenous covariates all indicate stationarity in the series' levels, which is consistent with the Brown and Yücel (2008) results, except for a small inconsistency associated with the natural-gas-storage variable. Brown and Yücel (2008) mention that “a non-stationary series means that any shock to the series will have permanent effects on it." One might expect shocks to deviations from five-year average natural-gas storage variables to have long-lasting effects. Consider a positive shock to natural-gas storage relative to the five-year average. In order for the difference between current storage levels and their 5-year average level to revert to its average, demand for natural gas would need to be above average subsequent to the positive storage shock, or production of natural gas would need to be below average, in order to induce increased storage withdrawals. This possibility of non-stationarity is manifested in the results of the ADF test on the STOR DIFF variable. The null hypothesis of non-stationarity is not rejected in the STOR DIFF variable unless a drift term is included in the underlying autoregression in which case the null hypothesis is rejected at the 1% level. The null hypothesis of non-stationarity is rejected at at least the 5% level for the levels of all of the exogenous covariates over the new full sample.

I ran additional stationarity tests to further investigate the nature of natural gas and crude oil price series, especially in light of the somewhat inconsistent results associated with the natural gas logged levels series. The second section in Table 1 shows the results of Phillips-Perron tests on gas and oil price series in logged levels and differences. The Phillips-Perron test is analogous to the ADF test, but applies Newey-West standard errors to account for autocorrelation rather than additional lags (Hamilton, 1994). The results of the Phillips-Perron tests are similar to ADF tests for both the replication and full samples, but the null hypothesis of non-stationarity is not rejected for the gas price series in any specification using the Phillips-Perron test.

Finally, I employed Zivot-Andrews unit root tests, which are analogous to the ADF tests, but allow for a single unknown structural break in the time series (Baum, 2015). The results of these tests are displayed in the third section of Table 1. The results of the Zivot-Andrews tests using the replicaton sample are consistent with the results from the other unit root tests. On the other hand, the Zivot-Andrews tests indicate strong evidence of stationarity in the logged natural gas price series using the full sample: this means that logged natural gas prices exhibit strong evidence of stationarity when a single structural break is controlled for. Interestingly, the t-statistic minimizing breakpoint (the break point that makes rejecting the null hypothesis of stationarity most likely) occurs in the summer of 2008 when the financial crisis was about to jolt commodity markets, and shale gas was beginning to impact the U.S. natural gas market.

The fourth section of Table 1 presents the results of Kwiatkowski-Phillips-Schmidt-Shin (KPSS) tests, which, in conjunction with unit root tests, might reveal fractional integration or a long-memory time-series process.[12] In contrast to the other unit root tests presented in Table 1, the null hypothesis of the KPSS tests is stationarity, which is rejected at at least the 5% level for all of the logged level price series in both the replication and full sample consistent with the majority of results from the other unit root tests. However, there is some evidence that logged levels of natural gas prices display long-memory stationarity: the null hypothesis of a unit root is rejected for logged natural gas prices when a trend is included in the ADF test using the replication sample, but the null hypothesis of no unit root is rejected for the same series in the KPSS test – together these imply the possibility of fractional integration (Lee and Schmidt, 1996); the same can be said for the logged natural gas series using the full sample, but without a trend.

The next step in a vector-error correction analysis of the effect of crude-oil prices on natural-gas prices is to apply the Johansen procedure to test for cointegration between the two series (Johansen, 1991 and Johansen, 1988). Tables 2 and 3 display the results of the Johansen tests for the existence of a cointegrating vector between crude-oil prices and natural-gas prices from the original Brown and Yücel (2008) analysis, and from my replication using the same sample period. The top halves of Tables 2 and 3 contain the results of the Johansen tests without including the exogenous covariates, while the bottom halves contain the results of these tests with the exogenous covariates.

Qualitatevely, the results are similar: the null hypothesis of zero cointegrating vectors between crude-oil prices and natural-gas prices is rejected with and without the inclusion of the exogenous covariates. Further, the rejection of the null hypothesis of zero cointegrating vectors becomes much stronger when the exogenous covariates are included. There does appear to be a small discrepancy in the magnitudes of the test statistics, however. I suspect this is because Brown and Yücel (2008) applied four lags in the underlying VAR when they performed the Johansen tests, because when I used four lags in the underlying VAR, my test statistics were almost exactly the same as the originals.[13] However, Table 3 in Brown and Yücel (2008), which is analogous to Tables 6 and 7 in this replication, includes four coefficients for lagged differences of logged crude-oil prices and logged natural-gas prices indicating that the underlying VAR contains five lags in the estimation of the VECM, as a lag is lost in the transformation.[14] Therefore, I have chosen to display results of the Johansen procedure with five lags rather than four. This would also explain another small discrepancy in the Brown and Yücel (2008) analysis. The normalized α coefficients in Table 2 of

[12]See Lee and Schmidt (1996) for the theoretical underpinnings related to KPSS and fractional integration, and see Zhang and Ji (2018) for an example of a long-memory time-series approach applied to oil and gas price cointegration. I thank an anonymous referee for pointing out this research.

[13]These results are available upon request, or can be generated easily by editing the code available with this article.

[14]See Hamilton (1994).

Brown and Yücel (2008) do not match the coefficients on the cointegrating terms in Table 3 of Brown and Yücel (2008), as they should, and do in this replication. This discrepancy in lags across the Johansen tests and VECM estimation would explain this difference.

Table 2: Bivariate Cointegration Tests (Henry Hub and West Texas Intermediate): Original Brown and Yücel (2008) Table

	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged 6/13/1997 to 6/8/2007 (without exogenous variables)	$p = 0$	0.0366	19.6577*	19.3116**
	$p \leq 1$	0.00067	0.3461	0.3461
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	-0.9322	
		0	(0.01083)	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0653	0.0084	
		(0.1083)	(0.0081)	
	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged 6/13/1997 to 6/8/2007 (with stationary exogenous variables)	$p = 0$	0.0786	43.472**	41.242**
	$p \leq 1$	0.0044	2.229	2.229
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	-0.8649	
		0	(0.0631)	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.1179	0.01626	
		(0.020)	(0.0107)	

Note: * and ** denote significance at better than 0.05 and 0.01 levels, respectively.

As shown by comparison of Table 2 Table 3, the β and α estimates from this replication are very similar to the analogous coefficients from the original Brown and Yücel (2008) analysis. All of the signs match, and the magnitudes are similar. The β estimates from both studies indicate that a 1% increase in the crude-oil price is associated with an approximately 0.9% increase in the natural-gas price in the long-run, while the α estimates indicate that natural-gas prices play the primary role in bringing prices back into alignment subsequent to deviations from the estimated long-run relationship.

Table 3: Bivariate Cointegration Tests (Henry Hub and West Texas Intermediate): Replication Table

	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged 6/13/1997 to 6/8/2007 (without exogenous variables)	$p=0$	**0.02998**	**16.1921***	**15.7342***
	$p \leq 1$	**0.00089**	**0.4579**	**0.4579**
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	**-0.9467**	
		0	**(0.1182)**	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0590	**0.0046**	
		(0.0156)	**(0.0080)**	
logged 6/13/1997 to 6/8/2007 (with stationary exogenous variables)	$p=0$	**0.0646**	**35.163****	**34.524****
	$p \leq 1$	**0.0012**	**0.6387**	**0.6387**
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	**-0.9118**	
		0	**(0.0680)**	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.1074	**0.0176**	
		(0.020)	**(0.011)**	

Note: * and ** denote significance at better than 0.05 and 0.01 levels, respectively.

Table 4 displays the results of the Johansen tests for the new full sample, while Table 4 displays the results the post sample period. These new results indicate that any cointegration that may have existed between logged crude-oil prices and logged natural-gas prices in the original sample is no longer present when the Johansen tests are applied to the full sample or the post sample. The Johansen test fails to reject the null hypothesis of zero cointegrating vectors over both samples whether or not the exogenous covariates are included. This represents an important change since the original Brown and Yücel (2008) analysis. The β estimates using the new full sample indicate that a 1% increase in crude-oil price is associated with 0.3% increase in natural-gas price in the long run, while the post-sample β estimate indicates that a 1% increase in crude-oil price is associated with a more than 1% increase in natural gas price in the long run. Comparison of the β estimates across the full and post samples makes clear that there has not been a clear-cut long-run stable relationship between U.S. crude-oil prices and U.S. natural-gas prices since at least the summer of 2008 – the same time when a structural break was found to occur in the logged natural gas price series in the Zivot-Andrews tests. Structural changes in natural-gas markets since 2008 appear to have broken the long-run relationship between crude-oil prices and natural-gas prices that Brown and Yücel (2008) (amongst others) discovered in 2008.

Table 4: Bivariate Cointegration Tests (Henry Hub and West Texas Intermediate): Full Sample Table

	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged 6/13/1997 to 6/16/2017 (without exogenous variables)	$p = 0$	0.000741	10.8701	7.7316
	$p \leq 1$	0.00301	3.1385	3.1385
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	-0.3108	
		0	(0.2411)	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0150	0.0010	
		(0.0055)	(0.0032)	
	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged 6/13/1997 to 6/16/2017 (with stationary exogenous variables)	$p = 0$	0.0680	10.5459	7.0968
	$p \leq 1$	0.0033	3.449	3.449
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	-0.4447	
		0	(0.2210)	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0122	0.0060	
		(0.0064)	(0.0038)	

Table 5: Bivariate Cointegration Tests (Henry Hub and West Texas Intermediate): Post Sample Table

	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged *6/15/2007* to 6/16/2017 (without exogenous variables)	$p = 0$	*0.0202*	*12.6657*	*10.6580*
	$p \leq 1$	*0.0038*	*2.0077*	*2.0077*
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	*-1.0520*	
		0	*(0.3337)*	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0150	0.0010	
		(0.0055)	(0.0032)	
	Ho: rank = p	Eigenvalue	Trace Statistic	Max Eigenvalue Statistic
logged *6/15/2007* to 6/16/2017 (with stationary exogenous variables)	$p = 0$	*0.0234*	*12.4583*	*12.3943*
	$p \leq 1$	*0.0001*	*0.0640*	*0.0640*
	Standardized Eigenvalues or βs with Standard Errors			
		Henry Hub	WTI	
		1	*-1.2818*	
		0	*(0.3337)*	
	Standardized α coefficients with Standard Errors			
		Henry Hub	WTI	
		-0.0098	*0.0161*	
		(0.0076)	*(0.052)*	

3.3 Model Results and Interpretation

Now, I turn to the results of the VECM estimation. Table 6 displays the VECM replication results in bold-face font next to the original Brown and Yücel (2008) results in normal font. Model 1 does not include the exogenous covariates, while Model 2 does. Table 7 displays the results of the analogous VECMs using the new full sample in normal font and the post sample in italic font. The estimated VECM coefficients using the replication sample are similar to the estimated coefficients presented in Brown and Yücel (2008). The only qualitative discrepancies that appear are sign differences between the coefficient on the third lag of crude-oil price in Model 1, the second lag of natural-gas price in Model 1, the first lag of natural-gas price in Model 2, and cooling degree days. None of these variables are individually statistically important in the original paper or this replication.

The coefficients on the lagged cointegrating errors displayed in the second row of Table 6, which match the replication α parameters in Table 2 and Table 3, are very similar across the original sample and the replication sample. These coefficients indicate that over the original sampling period, from June 1997 to June 2007, natural-gas prices adjusted when deviations from the long-run relationship with crude-oil prices occurred. Natural-gas price adjustment towards the long-run relationship with crude-oil prices occurred at a rate of approximately 5% per week in the model without the exogenous covariates, while gas prices adjusted at a rate of more than 10% per week towards the long-run relationship with crude-oil prices when the exogenous covariates are controlled for.

The statistically significant exogenous covariates all have their expected signs in both the original analysis and this replication. These coefficients are also quite similar in magnitude across the original analysis and replication. Deviations from normal heating degree days, e.g., cold surprises such as polar vortexes, led to increases in natural-gas prices over the original sample period. Similarly, deviations from normal cooling degree days led to higher natural-gas prices over the original sample. Lastly, increased differences between natural-gas storage levels and the five year average natural-gas storage level led to lower natural gas prices.

Now I turn to an examination of the VECM estimates associated with the full and post samples displayed in Table 7. These results should be interpreted with some care as the Johansen test failed to reject the null hypothesis of zero cointegrating vectors between natural-gas prices and crude-oil prices over both of these samples, as shown in Table 4 and Table 5, but I have chosen to present them for purposes of comparison with the original Brown and Yücel (2008) results.

In the model without the exogenous covariates, the coefficients on the lagged cointegrating terms are statistically significant at the 1% level for both the full and post samples. However, these coefficients are much smaller in magnitude than in the original and replication samples – natural-gas prices converge to their estimated long-run relationship with crude-oil prices at a rate of approximately 1.5% per week and 2% per week in the full and post samples, respectively. However, the Johansen procedure failed to reject the null hypothesis of zero cointegrating vectors between natural-gas prices and crude-oil prices, so this long-run relationship is not well defined in the data.

Table 6: Error-Correction Models of the Change in Natural Gas Price (original and **replication** samples)

	Model 1		Model 2	
constant	0.00028 (0.7545)	**0.0002** **(0.05)**	-0.0050 (-0.3794)	**-0.0017** **(-0.13)**
cointegrating term $(t-1)$	-0.0577 (-3.5274)**	**-0.0590** **(-3.78)****	-0.1158 (-5.7203)**	**-0.1074** **(-5.31)****
$\Delta P_{WTI}(t-1)$	0.1071 (1.2002)	**0.0686** **(0.79)**	0.0840 (0.9941)	**0.0647** **(0.77)**
$\Delta P_{WTI}(t-2)$	-0.1252 (-1.38900)	**-0.1241** **(-1.41)**	-0.0651 (-0.7633)	**-0.0967** **(-1.14)**
$\Delta P_{WTI}(t-3)$	0.0255 (0.2808)	**-0.0229** **(-0.26)**	0.0540 (0.6300)	**0.0025** **(0.03)**
$\Delta P_{WTI}(t-4)$	-0.0919 (-1.0215)	**-0.0598** **(-0.69)**	-0.1296 (-1.5181)	**-0.0968** **(-1.15)**
$\Delta P_{HH}(t-1)$	0.0925 (2.0693)*	**0.1501** **(3.37)****	-0.0301 (-0.6857)	**0.0987** **(2.22)***
$\Delta P_{HH}(t-2)$	0.0123 (0.2759)	**-0.0244** **(-0.54)**	-0.0336 (-0.7845)	**-0.0663** **(-1.49)**
$\Delta P_{HH}(t-3)$	-0.0351 (-0.7843)	**-0.0070** **(0.15)**	-0.0443 (-1.0172)	**-0.0034** **(-0.08)**
$\Delta P_{HH}(t-4)$	-0.0844 (-1.8811)*	**-0.0861** **(-1.93)**$^{+}$	-0.0602 (-1.4125)	**-0.0766** **(-1.77)**$^{+}$
HDD (t)			$2.08{\cdot}10^{-4}$ (2.4748)	$\mathbf{8.34{\cdot}10^{-5}}$ **(1.02)**
HDDDEV (t)			$1.06{\cdot}10^{-3}$ (5.6985)**	$\mathbf{9.11{\cdot}10^{-4}}$ **(4.84)****
CDD (t)			$-1.55{\cdot}10^{-4}$ (-0.6325)	$\mathbf{2.63{\cdot}10^{-4}}$ **(-1.10)**
CDDDEV (t)			$2.97{\cdot}10^{-3}$ (4.5411)**	$\mathbf{3.00{\cdot}10^{-3}}$ **(4.64)****
STORAGE DIFF (t)			$-5.69{\cdot}10^{-5}$ (-3.5139)**	$\mathbf{-4.39{\cdot}10^{-5}}$ **(-2.66)****
SHUT IN (t)			$1.16{\cdot}10^{-5}$ (2.3054)	$\mathbf{4.47{\cdot}10^{-4}}$ **(0.63)**
RMSE		**0.081**		**0.077**

Note: z-statistics are shown in parenthesis. $^{+}$,*, and ** denote significance at better than 0.1, 0.05 and 0.01 levels, respectively.

Table 7: Error-Correction Models of the Change in Natural Gas Price (full and *post* samples)

	Model 1		Model 2	
constant	0.00005 (0.02)	*-0.0002* *(-0.10)*	0.0009 (-0.09)	*0.0181* *(1.64)*
cointegrating term $(t-1)$	-0.01504 (-2.69)**	*-0.0211* *(-2.83)***	-0.0122 (-1.89)$^{+}$	*-0.0098* *(-1.28)*
$\Delta P_{WTI}(t-1)$	0.0474 (0.88)	*-0.0216* *(-0.33)*	0.0512 (0.95)	*-0.0425* *(-0.66)*
$\Delta P_{WTI}(t-2)$	-0.0721 (-1.33)	*-0.0504* *(-0.77)*	-0.0682 (-1.26)	*-0.0725* *(-1.11)*
$\Delta P_{WTI}(t-3)$	-0.0156 (-0.29)	*-0.0538* *(-0.83)*	-0.0047 (-0.09)	*-0.0786* *(-1.21)*
$\Delta P_{WTI}(t-4)$	-0.0369 (-0.69)	*-0.0602* *(-0.69)*	-0.0511 (-0.96)	*-0.0882* *(-1.37)*
$\Delta P_{HH}(t-1)$	0.1330 (4.24)**	*0.1421* *(3.23)***	0.0974 (3.03)**	*0.1106* *(2.41)**
$\Delta P_{HH}(t-2)$	-0.0261 (-0.82)	*0.0015* *(0.03)*	-0.0590 (-1.82)$^{+}$	*-0.0156* *(-0.34)*
$\Delta P_{HH}(t-3)$	-0.0520 (-1.64)	*-0.1233* *(-2.77)***	-0.0665 (-2.10)*	*-0.1362* *(-3.05)***
$\Delta P_{HH}(t-4)$	-0.0481 (-1.53)	*0.0504* *(1.13)*	-0.0564 (-1.80)$^{+}$	*0.0379* *(0.84)*
HDD (t)			$-5.77 \cdot 10^{-5}$ (0.95)	*$-1.27 \cdot 10^{-4}$* *(-1.88)*$^{+}$
HDDDEV (t)			$5.06 \cdot 10^{-4}$ (3.91)**	*$1.51 \cdot 10^{-4}$* *(0.88)*
CDD (t)			$-3.98 \cdot 10^{-4}$ (-2.52)*	*$-5.01 \cdot 10^{-4}$* *(-2.52)**
CDDDEV (t)			$1.31 \cdot 10^{-3}$ (3.20)**	*$4.34 \cdot 10^{-4}$* *(0.86)*
STORAGE DIFF (t)			$1.15 \cdot 10^{-5}$ (1.47)**	*$1.83 \cdot 10^{-5}$* *(2.25)**
SHUT IN (t)			$-1.69 \cdot 10^{-4}$ (-0.26)	*-0.0396* *(-2.01)**
RMSE	0.073	*0.063*	0.072	*0.062*

Note: z-statistics are shown in parenthesis. $^{+}$,*, and ** denote significance at better than 0.1, 0.05 and 0.01 levels, respectively.

Interestingly, when the exogenous covariates are included in the model, both the magnitude and the statistical significance of the coefficients on the lagged cointegrating terms decreases, which contrasts the analogous coefficients using the original and replication samples. This coefficient is significant at the 10% level for the full sample and implies that natural-gas prices converge to the long-run relationship with crude-oil prices at a rate of 1.2% per week, but again no statistically significant long-run relationship between natural-gas prices and crude-oil prices even exists over this sample. The analogous coefficient for the post sample is not statistically significant. The results of the VECM estimation over the full and post samples presented in Table 7 in conjunction with the results of the Johansen procedure over the full and post samples presented in Tables 4 and 5 make clear that any long-run relationship between natural-gas and crude-oil prices that existed over the original and replication samples has since been broken – likely as a result of increased natural-gas supply associated with the shale revolution since 2008.

Important differences between the estimated coefficients on the exogenous covariates between the original sample period and the full and post sample periods are also apparent from inspection of Tables 6 and 7. Deviations from normal heating degree days and deviations from cooling degree days both had statistically significant positive impacts on natural-gas prices over the full sample period, but the magnitude of these effects is much smaller over the full sample than over the original sample. Further, this statistical significance disappears when the model is estimated over the post sample. This plausibly results from the increased supply of natural gas available during the post-sample period: cold weather surprises have had less of an impact on natural-gas prices recently, as future shortages resulting from current extreme weather events became less likely. Surprisingly, the coefficient on the STORAGE DIFF variable is positive and statistically significant in both the full and post samples, as opposed to being negative and statistically significant in the original and replication samples.

Finally, I turn to a visual examination of the fit of the VECM model that includes the exogenous covariates (Model 2) over the full sample. Figure 3 displays the actual Henry Hub natural-gas price and the fitted values from Model 2 over the full sample in its top panel, and the residuals in its bottom panel. Figure 3 is analogous to Figure 3 in Brown and Yücel (2008). Although the VECM results over the full and post samples suggest that a long-run relationship between natural-gas prices and crude-oil prices has not persisted during the last decade, Figure 3 shows that the VECM model that includes exogenous covariates explains natural gas prices very well relative to simple rules of thumb and burner-tip parity rules. In fact, the VECM appears to explain natural gas prices after 2008 with even more accuracy than it does before 2008, when the long-run relationship with oil was present, as the maximum magnitude of the residuals is less than $2.00 in the post sample period, while the maximum residual in the over the original sample period is more than $5.00. Figure 3 makes clear that the VECM model continues to be useful in answering: what drives natural gas prices?

Evidence of a better fit is buttressed by examination of the RMSE statistics for each sample period, which are displayed in the last rows of Tables 6 and 7: the RMSEs are largest in the original sample period, before 2008, and are smallest in the post sample period, after 2008. Further, the exogenous covariates have much less of an impact on the RMSEs in the full and post sample models. Comparison of coefficients across Tables 6 and 7 indicate that lagged natural gas prices have come to be the most important predictor of current natural gas prices since 2008, which again dovetails with our earlier fining that natural gas prices have been stationary conditional on a structural break that occurred around the summer of 2008.

Figure 3: Actual and Fitted Natural Gas Prices (Model 2, Full Sample)

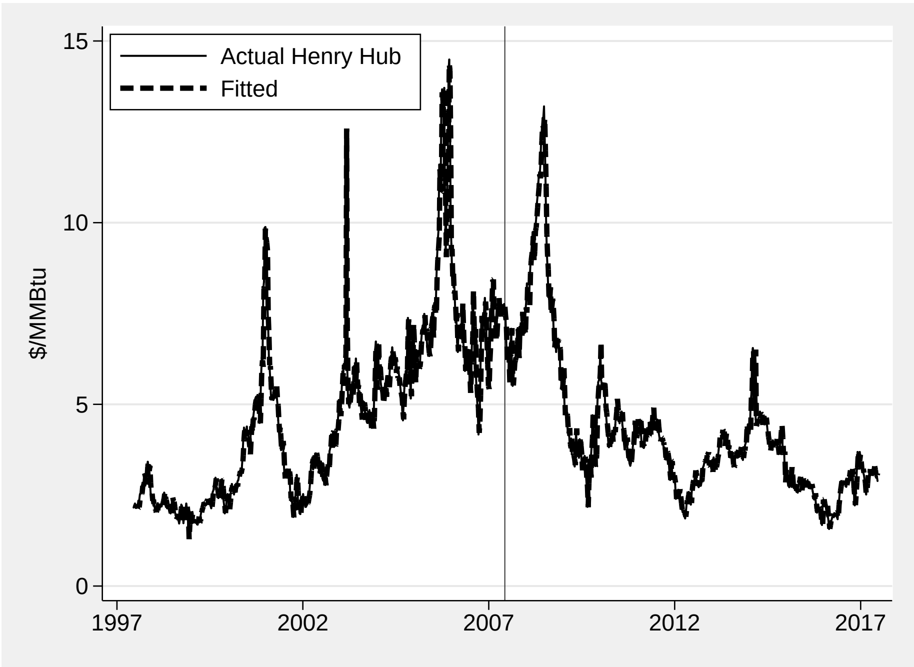

(a) Actual and Fitted Natural Gas Prices

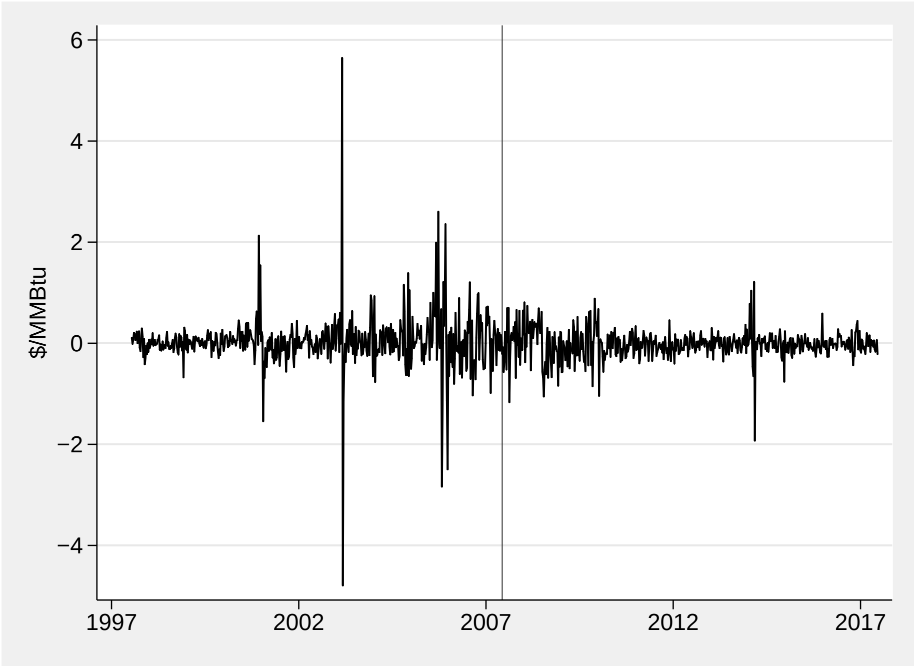

(b) Differences between the Actual and Fitted Natural Gas Prices

4 Conclusion

This replication of Brown and Yücel (2008) confirms those authors' original results that indicated a long-run relationship between natural-gas prices and crude-oil prices existed between June 1997 and June 2007. However, when the sample period is updated to include weekly data until June 2017, this long-run relationship appears to have broken: the Johansen procedure fails to reject the null hypothesis of zero cointegrating vectors between natural-gas prices and crude-oil prices when the procedure is applied to the full new sample (June 1997 to June 2017) and when the procedure is applied to the post sample (June 2007 to June 2017). On the other hand, a model that includes the exogenous covariates that were included in the original Brown and Yücel (2008) analysis continues to explain natural-gas prices quite well.

The same year that Brown and Yücel (2008) published their highly influential analysis of the drivers of natural gas prices the global financial crisis caused a precipitous fall in crude oil and natural gas prices along with the prices of other commodities, as shown Figure 1. Crude oil prices recovered relatively quickly, but natural gas prices remained depressed, and the historical relationship between crude oil prices and natural gas prices appeared to break. This replication analysis reaffirms the relationship that Brown and Yücel (2008) discovered, but confirms that the relationship broke soon after that stusy was published. The financial crisis was likely the original cause of the initial fall in natural gas prices, but a longer lasting trend appears to have been at play that appears to have caused a structural shift in the trajectory of natural gas prices.

The most likely reason that the long-run relationship between natural-gas prices and crude-oil prices broke after June 2007 was the onset of the "shale-gas revolution." Natural gas supply increased dramatically between 2005 and the present in the United States, as natural gas producers unlocked massive quantities of gas from shale reservoirs by applying new technology. This new supply led to dramatically lower natural gas prices, while crude-oil prices remained relatively high until 2014. However, similar technologies are currently dramatically increasing the supply of crude oil in the U.S. The expanded natural gas supply in the U.S. is creating increased demand for natural gas, as more natural gas electricity generation capacity is added, liquified natural gas export facilities are built, and the petrochemical industries in the U.S. are expanded. These and other factors will continue to make natural-gas and crude-oil markets ripe for economic analysis for years to come.

References

Baum, C. (2015). “Zandrews: Stata module to calculate zivot-andrews unit root test in presence of structural break.” RePEc:boc:bocode:s437301.

Brigida, M. (2014). “The switching relationship between natural gas and crude oil prices.” *Energy Economics* 43, 48–55.

Brown, S. P. A., S. A. Gabriel, and R. Egging (2010). “Abundant shale gas resources: Some implications for energy policy.” Resources for the Future. www.rff.org/documents/195/RFF-BCK-Brownetal-ShaleGas.pdf.

Brown, S. P. A. and M. K. Yücel (2008). “What drives natural gas prices?” *The Energy Journal* 29, 45–60.

Caporin, M. and F. Fontini (2017). “The long-run oil–natural gas price relationship and the shale gas revolution.” *Energy Economics* 64, 511 – 519.

EIA (2016). *Annual energy outlook 2016.* Technical report, U.S. Energy Information Administration.

Erdős, P. (2012). “Have oil and gas prices got separated?” *Energy Policy* 49, 707–718.

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

Hartley, P. R., I. Medlock, Kenneth B., and J. E. Rosthal (2008). “The relationship of natural gas to oil prices.” *The Energy Journal* 29, 47-66. DOI: 10.5547/ISSN0195-6574-EJ-Vol29-No3-3.

Johansen, S. (1988). “Statistical analysis of cointegration vectors.” *Journal of Economic Dynamics and Control* 12, 231–254.

Johansen, S. (1991). “Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models.” *Econometrica* 59, 1986 – 1998.

Lee, D. and P. Schmidt (1996). “On the power of the kpss test of stationarity against fractionally-integrated alternatives.” *Journal of Econometrics* 73, 285–302.

Liu, H. and J. Li (2018). “The US shale gas revolution and its externality on crude oil prices: A counterfactual analysis.” *Sustainability* 10, 1-17.

Nick, S. and S. Thoenes (2014). “What drives natural gas prices? —a structural var approach.” *Energy Economics* 45, 517–527. DOI: 10.1016/j.eneco.2014.08.010.

Pettersson, F., P. Söderholm, and R. Lundmark (2012). “Fuel switching and climate and energy policies in the european power generation sector: A generalized leontief model.” *Energy Economics* 34, 1064–1073. DOI: doi.org/10.1016/j.eneco.2011.09.001.

Ramberg, D. J. and J. E. Parsons (2012). "The weak tie between natural gas and oil prices." *The Energy Journal* 33, 13–35.

Zhang, D. and Q. Ji (2018). "Further evidence on the debate of oil-gas price decoupling: A long memory approach." *Energy Policy* 113, 68–75.