WORKING PAPER SERIES

Natural Gas and U.S. Industrial Production: A Closer Look at Four Industries

Vipin Arora and
Elizabeth Sendich

August 30, 2014

Natural Gas and U.S. Industrial Production: A Closer Look at Four Industries *†

Abstract

We consider the relationship between natural gas prices and production in the U.S. resins, agricultural chemicals, cement, and aluminium industries. Overall, our analysis using various tests and regressions does not allow for generalizations about the association between natural gas prices and production across these four energy-intensive industries. Rather, the relationships between natural gas and production appear to be driven by the particular institutional details of each industry.

Introduction

How important are natural gas prices to production in U.S. energy-intensive industries? How beneficial have the recent increases in American natural gas supply been to these industries? If approved, will exports of U.S. liquefied natural gas (LNG) negatively impact the industries which use this fuel? These are important questions that have received much attention since the beginning of the U.S. shale gas revolution. But there is little empirical work on which to base any answers.

In this paper we consider the importance of natural gas prices for output in four specific U.S. energy-intensive industries: resins, agricultural chemicals, cement, and aluminum. Our primary goal is to shed some light on the questions above using data that includes the time period after shale gas became an important component of U.S. natural gas supply. We also seek to understand and account for the results in terms of the key features specific to each of these industries.

---

*The analysis and conclusions expressed here are those of the authors and not necessarily those of the U.S. Energy Information Administration.

†Katherine Calais provided excellent research assistance. We have also benefitted from the comments and suggestions of Joe Benneche, Stephen Brown, John Conti, Peter Gross, Fred Joutz, Thomas Lee, Kay Smith, Russell Tarver, and members of the American Statistical Association Advisory Committee on Energy Statistics.

Many recent studies use both macroeconomic and input-output models to evaluate the economic importance of natural gas for specific U.S. industries, but empirical work is limited.[1] In one example, Kliesen (2006) finds that natural gas prices have historically been unable to predict industrial production in U.S. manufacturing industries. This paper was completed before the large increases in natural gas production due to shale gas. In additional research completed before this time, Costello et al. (2006) analyze the interaction between industrial natural gas prices, natural gas consumption, and industrial sector activity. They conclude that industrial sector firms respond to relative energy prices, including natural gas prices. Weber (2012) considers the economic impacts of the natural gas market after the shale gas boom, but on a regional level.[2] Arora and Lieskovsky (forthcoming) also consider the period after 2008, although they focus on aggregate impacts of the natural gas market.

Resins, agricultural chemicals, cement, and aluminum are chosen in the analysis because each is an energy-intensive industry that vary in how much they use natural gas. Both resins and agricultural chemicals are heavy users of natural gas. Aluminum is more of an intermediate natural gas user, both by volume and as a percentage of total fuel use/mix. Cement uses very little natural gas volume relative to the other three, and has a broad fuel mix. By studying these four industries we are able to get an overview of how different parts of the manufacturing sector may respond to natural gas prices.

With these various considerations in mind, we begin our analysis by conducting a variety of different Granger causality tests of natural gas prices on each of the four industrial production indices. For each industrial production series we vary the tests by changing the natural gas price, altering lag lengths, and modifying the specifications to include the real oil price and total industrial production. This results in a total of 144 tests each for resins, agricultural chemicals, cement, and aluminum. In general the results are inconclusive, and with the exception of resins only a small fraction of tests are consistent with the ability of natural gas prices to predict industrial production: 25% for resins, 10.4% for agricultural chemicals, 7.6% for cement, and 12.5% for aluminum.

The factors behind the inconclusiveness of these results may be due to well-known limitations of Granger causality tests, namely the omission of important variables and uncertainty over the proper lag length. However, in this case they may also be due to the assumptions of a linear relationship between natural gas prices and industrial production or that the estimated relationships between these variables have been constant over time. The latter explanation is a particularly likely culprit in light of recent events in U.S. natural gas markets.

---

[1] Examples of studies focused on the economy-wide and industry-specific impacts of natural gas include ACC (2011), ACC (2012), Arora (2013), CitiGPS (2012), CRA (2013), ICF (2012), IHS (2011), IHS (2012), IHS (2013), McKinsey Global Institute (2013), NERA (2012), Peterson Institute (2014), PWC (2011), PWC (2012b), PWC (2012a), PWC (2013), and Kinnaman (2010) and citations therein.

[2] In research before deregulation of the U.S. natural gas market, Leone (1982) claims the impact of natural gas price increases on the northeast regional economy are at least offset by gains due to greater revenue for natural gas producers.

These limitations lead us to use non-parametric methods, specifically locally weighted regression (LOESS) to understand the relationship between natural gas prices and production in our chosen industries. Our primary specifications have three variables (industrial production, natural gas prices, and total industrial production), and consider three and six month lags of natural gas prices. The results from regressing each specific industrial production index on the other two variables vary between industries. Both resins and aluminum show an inverse association between natural gas prices and production, while the results for agricultural chemicals and cement indicate an association in the same direction. Overall, neither the Granger causality tests nor the non-parametric regressions allow for generalizations about the relationship between natural gas prices and production in these four energy-intensive industries.

Data and descriptive statistics

This section outlines the data series used in subsequent analysis, including the sources, time periods under consideration, and any modifications. It then describes the modified series in more detail, and presents correlations between industrial production and natural gas prices.

Data

We use nine different monthly series in our analysis. There are four industry-specific production indices and a total U.S. industrial production index. The four specific indices correspond to resins [NAICS 3252], agricultural chemicals [NAICS 3253], cement [NAICS 3273], and aluminum [NAICS 3313] sectors. The total industrial production index includes the U.S. manufacturing sector as well mining and utilities. Each of the industrial production indices are taken from the Board of Governors of the Federal Reserve over the period 1990M01-2013M04, and are seasonally unadjusted.[3]

Three different natural gas prices and one oil price are also utilized. The nominal price of natural gas at Henry Hub is taken from the Federal Reserve Economic Data (FRED) system at the Federal Reserve Bank of St. Louis from 1993M11-2013M04.[4] We deflate this by the seasonally unadjusted Producer Price Index (PPI) for fuels [WPU05], which comes from the U.S. Bureau of Labor Statistics (BLS).[5] The seasonally unadjusted nominal U.S. industrial price for natural gas is obtained from the U.S. Energy Information Administration (EIA) for the period 2001M01-2013M04.[6] This price and the composite refiner acquisition costs of crude oil (also from the EIA) are both deflated in the same manner as the Henry Hub price. The final natural gas price we use is the PPI for the industrial natural

[3]See http://www.federalreserve.gov/releases/g17/download.htm.
[4]See http://research.stlouisfed.org/fred2/series/GASPRICE/.
[5]See http://data.bls.gov/timeseries/WPU05.
[6]See http://www.eia.gov/dnav/ng/ng_pri_sum_a_EPG0_PIN_DMcf_m.htm.

gas price [WPU0553] from the BLS, which is available from 1990M01-2013M04.[7]

Each of the nine series appear to be trending over the sample period.[8] This sample period is 1990M01-2013M04 for the industrial production indices, PPI natural gas price, and real refiner acquisition costs of crude oil; 1993M11-2013M104 for the real Henry Hub price; and 2001M01-2013M04 for the real industrial natural gas price. Each of the series save the real oil price also have seasonal movements. In our analysis we work with the annual percent change in each variable to remove both the trend and seasonality in the series.

Descriptive statistics

The annual percent changes in the specific industrial production series and each natural gas price are plotted in Figure 1. Both sets of series move together in general, but have pronounced differences. Especially notable is that before 2003, there are large changes in the real Henry Hub price which are not necessarily reflected in either the real industrial or PPI industrial prices. But the price series tend to move similarly after 2003. In terms of industrial production indices, before 2009 the series move in a similar fashion. After this time resins and agricultural chemicals show large variations, and aluminium and agricultural chemicals production have movements which are relatively synchronized.

Some of these differences and similarities are reflected in the descriptive statistics in panels (a) and (b) of Table 1. Panel (a) shows that the average growth rate in resins production ($\%\Delta_{RS}$) has a mean value close to zero over the sample, and a standard deviation of 9.2%. Resins also has strong negative skewness. The kurtosis for agricultural chemicals production ($\%\Delta_{AC}$) shows that the observed values of changes in this variable have a distribution which differs from the normal (kurtosis = 0), at 12.39. Changes in this series have positive skew, indicating that values fall below the mean more often than above. This mean value is slightly negative, at -0.4%, and the standard deviation is just over 8%.

The annual percent changes in cement production ($\%\Delta_{CM}$) are also close to zero, and the series is as volatile as resins. Changes in cement are also highly negatively skewed, implying the values in the sample tend to fall above the mean, with fewer below. And the kurtosis value of 7.50 shows that the observed values of changes in organic chemicals production have a distribution which also differs from the normal. Changes in aluminum production ($\%\Delta_{AL}$) have the lowest standard deviation of the three series, at 1% over the sample period. The mean is a slightly positive 0.6%, with positive skewness of 0.166 and kurtosis similar to organic chemicals at 4.83.

Panel (b) of Table 1 shows that changes in the natural gas price series are much more volatile than the corresponding changes in the industrial production indices. Consistent with Figure 1, changes in real

[7] See http://data.bls.gov/timeseries/WPU0553?include_graphs=false&output_type=column&years_option=all_years.

[8] In each case unit root tests indicate I(1) series: the Augmented Dickey-Fuller (ADF) test cannot reject the null of a unit root, and the KPSS test rejects the null of stationarity.

Henry Hub prices ($\%\Delta_{HH}$) are the most volatile, with a standard deviation of over 31%, the PPI for industrial natural gas ($\%\Delta_{PP}$) has a standard deviation of 21%, while changes in the real industrial natural gas price ($\%\Delta_{IN}$) have a standard deviation just below 19%.

The mean values of $\%\Delta_{HH}$ and $\%\Delta_{IN}$ are negative over their respective samples (-3.6% and -6.2%), while the mean value of $\%\Delta_{PP}$ is positive (1.9%). All three natural gas price series are positively skewed, and their kurtosis values indicate that the observed values have a narrower distribution than the normal. In each of the price statistics, changes in the real Henry Hub price stand out from the other two series.

Panel (c) of Table 1 begins to explore the relationship between changes in each industrial production index and each natural gas price through non-parametric correlations. We use non-parametric correlations because of the large changes which have occurred in U.S. natural gas markets over the past 25 years. Each row shows the Spearman correlation coefficient for changes in the respective industrial production index and changes in the three natural gas prices at lags of three and six months.[9] These particular lags account for the fact that changes in natural gas prices take time to have an impact on industrial production.

The results indicate that there is a negative relationship between growth in resins or aluminum production and growth in any of the natural gas prices. These are all statistically significant and hold at either lag of the natural gas price (with the exception of resins and aluminum production and three lags of the real Henry Hub price). Each of these two industries clearly have an association with past changes in natural gas prices. The same is not true for either cement or agricultural chemicals production. The correlations with cement are positive, some of which are statistically significant, while agricultural chemicals production does not have a statistically significant correlation with any of the lagged natural gas prices.

Quantitative analysis

In this section we further explore the relationship between changes in natural gas prices and changes in U.S. resins, agricultural chemicals, cement, and aluminum production. The section begins with standard Granger causality tests of each industrial production index on each natural gas price. Concerns about the stability of each price series, as well the assumption of linearity inherent in the tests leads us to move to non-parametric methods. The results of non-parametric regressions of the

[9]The Spearman correlation coefficient is calculated using the same formula as the standard Pearson correlation coefficient. The difference is that it is applied to ranked data and not each industrial production/natural gas price pair. To construct the correlation each of the observed values for changes in a specific industrial production index are arranged in ascending order, as are the observed values for a specific natural gas price. These ordered values are paired, and then used in calculation of the correlation coefficient using the standard Pearson formula.

industrial production indices on the natural gas prices concludes the section.

Granger causality tests

Granger causality tests are a standard method used to explore the ability of past values in one series to predict another. Although we do not interpret these tests as inferring causation from one series to another, they do provide important information on the association between changes in natural gas prices and changes in industrial production indices. We use several specifications of these tests to account for other possible factors and estimation issues as well [see for example the discussions in Brown and Yucel (2008) and Joutz and Villar (2006)].

Technical details and specifications

We begin by conducting an analysis similar to Kliesen (2006) and attempt to quantify the importance of changes in real natural gas prices for the respective industrial production indices. The methods we use are based on the literature on oil price shocks and economic activity [see Hamilton (1996) and Hooker (1996) for notable examples], specifically the use of bivariate Granger causality tests as in Cunado and Perez de Gracia (2003). We also extend the tests to three and four variables to control for the impact of total industrial production and the real oil price.

Our initial bivariate specification tests for Granger causality from past annual percent changes in each natural gas price (Δ_{xx}, where $xx = HH$, IN, or PP) to the annual percent change of each industrial production index (Δ_{yy}, where $yy = OC$, AC, or AL) are specified as:

$$\%\Delta_{yy,t} = \alpha_0 + \Sigma_{j=1}^{k}\alpha_{1j}\%\Delta_{yy,t-j} + \Sigma_{i=1}^{k}\alpha_{2i}\%\Delta_{xx,t-i} + \epsilon_t \tag{1}$$

where ϵ_t is a normally distributed error term with an expected value of zero. We test the hypothesis that $\alpha_{21} = \alpha_{22} = ... = \alpha_{2k} = 0$. Rejecting this hypothesis indicates that changes in the respective natural gas price Granger cause changes in the particular industrial production index. The lag length (k) is set at $k = 3$, 6, and 12 months.

For the single-equation bivariate case specified above, each index is regressed against the three natural gas prices at three different lags, for a total of 9 specifications. The sample is also varied, from the beginning of the price data to 2013M04, and from the beginning of the price data through 2007M12. This gives a total of 18 Granger causality tests for the single-equation bivariate case of each industrial production index. These same tests are also repeated with the equation as part of a vector auto-regressive (VAR) system of equations, for a total of 36 Granger causality tests of each industrial production index in the two variable specification.

We then extend the two-variable case to consider the annual percent changes in real oil prices

($\%\Delta_{rpo}$) as well:

$$\%\Delta_{yy,t} = \alpha_0 + \Sigma_{j=1}^{k}\alpha_{1j}\%\Delta_{yy,t-j} + \Sigma_{i=1}^{k}\alpha_{2i}\%\Delta_{xx,t-i} + \Sigma_{l=1}^{k}\alpha_{3l}\%\Delta_{rpo,t-l} + \epsilon_t \quad (2)$$

Again, we test the hypothesis that $\alpha_{21} = \alpha_{22} = ... = \alpha_{2k} = 0$. As in the two variable case these are estimated over three different lags, two different sample periods, and as single-equations or VARs for each industrial production index. This adds another 36 Granger causality tests to give a total of 72 for each industrial production index over the two and three variable specifications.

We extend the two-variable case again to consider the annual percent changes in total industrial production ($\%\Delta_{ip}$):

$$\%\Delta_{yy,t} = \alpha_0 + \Sigma_{j=1}^{k}\alpha_{1j}\%\Delta_{yy,t-j} + \Sigma_{i=1}^{k}\alpha_{2i}\%\Delta_{xx,t-i} + \Sigma_{l=1}^{k}\alpha_{3l}\%\Delta_{ip,t-l} + \epsilon_t \quad (3)$$

and also to four variables:

$$\%\Delta_{yy,t} = \alpha_0 + \Sigma_{j=1}^{k}\alpha_{1j}\%\Delta_{yy,t-j} + \Sigma_{i=1}^{k}\alpha_{2i}\%\Delta_{xx,t-i} + \Sigma_{l=1}^{k}\alpha_{3l}\%\Delta_{ip,t-l} + \Sigma_{u=1}^{k}\alpha_{4u}\%\Delta_{rpo,t-u} + \epsilon_t \quad (4)$$

This gives a total of four specifications, each with 36 different tests, for a total of 144 Granger causality tests of each industrial production index.

Granger causality test results

Instead of listing the results for each individual Granger causality test (a total of 576), Table 2 summarizes them by industry. The first column specifies the number of tests for each industry: there are four different specifications for each industry, and each specification has 36 different Granger causality tests due to variations in lag length, natural gas price, sample period, and estimation method. The second column lists the number and fraction of the tests where the null of no Granger causality can be rejected. The final column lists the number and fraction of those tests where the null is rejected and the variables are not jointly determined.[10]

Table 2 shows that there is no conclusive evidence one way or another about whether changes in natural gas prices can predict changes in any of the industrial production indices. Although a relatively low percentage of the tests are able to reject the no Granger causality null in any of the industries (with the exception of resins), there are some tests which do reject this null. In cement, slightly over

[10] That is, the null that the natural gas price does not Granger cause industrial production is rejected, but the null that the industrial production index does not Granger cause the natural gas price cannot be rejected.

7.5% of tests indicate this predictive power. This rises to 10.4% for agricultural chemicals, 12.5% for aluminum, and 25% for resins.

The results are inconclusive because, as with any Granger causality analysis, there is uncertainty over the correct specification and number of lags. The results in Table 2 just highlight this uncertainty in the case of these four specific industries and natural gas prices. In this context there are also other concerns with the tests that are detailed in the next sub-section.

Non-parametric regressions

The applicability of the various Granger causality tests outlined above in understanding the relationship between industrial production and natural gas prices is questionable for reasons other than lags and specification as well. First, there is an assumption of a linear relationship between the variables which may not be the case. Additionally, the specification also implies a symmetric relationship; but there is no reason to exclude the possibility that industrial production responds differently when changes in natural gas prices are positive rather than negative. Finally, estimation by OLS assumes that the coefficients in the Granger causality test equations are constant over time. In light of recent events in U.S. natural gas markets, this seems unlikely [see Arora and Lieskovsky (forthcoming)]. These limitations lead us to use non-parametric methods, specifically locally weighted regressions (LOESS) [see Cleveland (1979)].

Technical details and specifications

The LOESS method is a type of non-parametric regression that does not require assumptions about the functional form of the relationship between natural gas prices and industrial production, nor does it force us to assume constant parameter values. To illustrate, consider the general relationship between one of the industrial production indices and natural gas prices:

$$\%\Delta_{yy,t} = f(\%\Delta_{xx,t}) + \epsilon_t \tag{5}$$

where ϵ_t is a normally distributed error term with an expected value of zero. The Granger causality tests outlined above assume that $f()$ is linear: $= \alpha_0 + \Sigma_{j=1}^{k}\alpha_{1j}\%\Delta_{yy,t-j} + \Sigma_{i=1}^{k}\alpha_{2i}\%\Delta_{xx,t-i}$. And estimation of the parameters (α) by OLS assumes that they do not change over time. In contrast, a non-parametric regression estimates $f()$ directly without imposing these restrictions. In the case of LOESS this is done by performing a regression at each particular value of the natural gas price.

That is, LOESS provides an estimate of $\%\Delta_{yy,t}$ ($\%\widehat{\Delta}_{yy,t}$) by performing a regression at each value of $\%\Delta_{xx,t}$. This is called a local (or nearest neighbor) regression because the user specifies a bandwidth parameter that defines the neighborhood around $\%\Delta_{xx,t}$ which is used for each regression. For example, suppose there are 10 observations of industrial production and natural gas prices in our sample. To estimate $f()$ using LOESS, one needs to pick each individual $\%\Delta_{xx,t}$ and construct a regression at that point. The local area for the regression is defined by choosing the bandwidth, which is the fraction of other $\%\Delta_{xx,t}$ values to be included in the sample. If the user picks a bandwidth of 0.5, then the five points closest to each $\%\Delta_{xx,t}$ are used in the local regression.

LOESS is a weighted regression because each of the five points chosen for the local regression in the example above will be weighted differently when constructing $\%\widehat{\Delta}_{yy,t}$. Various weighting functions can be used, but the main point is that of the five points included in the local regression, those further away from $\%\Delta_{xx,t}$ will be assigned a lower weight, and thus will have less of an impact on $\%\widehat{\Delta}_{yy,t}$ than values which are closer. In addition, each locally weighted regression can be linear or quadratic (in which case the procedure is known as LOWESS). This basic two variable example can be extended to three and even four variables, although the computations take substantially longer as variables are added. One can also use lags of either variable instead of the current values shown above.

In our analysis we use three variable LOESS regressions of the industrial production indices on either three or six lags of the natural gas prices and one lag of total industrial production (to control for demand). As with the correlations, the lags account for the fact that it takes time for changes in natural gas prices to impact industrial production. In terms of the model, this can be written:

$$\%\Delta_{yy,t} = f(\%\Delta_{xx,t-j}, \%\Delta_{ip,t-1}) + \epsilon_t \quad j = 3, 6 \tag{6}$$

where the bandwidth is either 0.3 or 0.5, each local regression is linear, and a tri-cube function is used in weighting. Although the selection of the lag lengths, bandwidth, and order of the linear regressions seem arbitrary, we experimented with many other combinations. Results are available for combinations of the regressions where there are two variables, the bandwidth ranges from 0.1 to 0.9, the local regressions are quadratic, and the lags range from 0 to 12. The results presented here are representative of these other variations.

Results

The non-parametric regression results are shown in Figures (2) - (5). Each column of the figures uses a different natural gas price, while the top row uses a lag length of three for the price and the bottom row a lag length of six. The solid line in each figure is a regression with a bandwidth of 0.3 and the dashed line uses a bandwidth of 0.5.

In the case of the resins and aluminum industries, Figures (2) and (5) have a straight forward interpretation, as with the correlations discussed above. Across each natural gas price, both regressions show increasing production in these industries is associated with falling prices, as well as an association between decreasing production and rising prices. The trends from the regressions on the resins industry most likely reflect natural gas price effects in dry natural gas fuel and liquid co-products (which are used as feedstock). Aluminum regressions point to natural gas fuel use and some price effects from electricity use as natural gas begins to make inroads in the generation fuel mix. Because the resins industry is concentrated in the southeastern U.S, and closely tied to the petrochemicals industry, the PPI results are the most relevant. For the aluminum industry any price is applicable.

Figures (3) and (4) show that the agricultural chemicals and cement industries have results which differ from those for resins and aluminum. Because the cement industry is composed of multiple small establishments that are regionally dispersed, the Henry Hub price is most relevant for this particular analysis. When analyzing the agricultural chemicals industry it is important to recognize the diversity of sub-industry composition. The primary sub-industries are: nitrogenous fertilizers, phosphatic fertilizers, all other fertilizers, and herbicides and pesticides. The disparate nature of these various industries makes the Henry Hub price results most relevant.

Still, across prices both industries have results that indicate increasing production is associated with rising prices and decreasing production is associated with falling prices. Understanding the reason for the price change is very important for agricultural chemicals, as its products go into critical inputs for other industries, primarily construction. The results might then reflect the fact that an increased natural gas price stimulates natural gas production, which can drive construction. However, one could also propose a scenario where the construction sector is growing due to broader economic growth, and this broader growth is raising demand for natural gas alongside the price. The same is true for cement, although the inverse association is much less pronounced in this case.

Overall, our regressions do not allow for generalizations about the association between natural gas prices and production across these four energy-intensive industries. Rather, the relationships between natural gas and production appear to be driven by the particular institutional details of each industry.

Conclusion

In this paper we consider the importance of natural gas prices for output in four specific U.S. energy-intensive industries: resins, agricultural chemicals, cement, and aluminum. These four are chosen in the analysis because each is an energy-intensive industry that varies in how much they use natural gas. We begin our analysis by conducting a variety of different Granger causality tests of natural gas prices on each of the four industrial production indices. In general the results are inconclusive, and with the exception of resins only a small fraction of tests are consistent with the ability of natural gas prices to predict industrial production: 25% for resins, 10.4% for agricultural chemicals, 7.6% for cement, and 12.5% for aluminum.

These limitations lead us to use non-parametric methods, specifically locally weighted regression (LOESS) to understand the relationship between natural gas prices and production in our chosen industries. The results from regressing each specific industrial production index on a natural gas price and total industrial production vary between industries. Both resins and aluminum show an inverse association between natural gas prices and production, while the results for agricultural chemicals and cement indicate an association in the same direction. Overall, neither the Granger causality tests nor the non-parametric regressions allow for generalizations about the relationship between natural gas prices and production in these four energy-intensive industries.

References

ACC, "Shale Gas and New Petrochemicals Investment: Benefits for the Economy, Jobs, and US Manufacturing," 2011.

_ , "Shale Gas, Competitiveness and New U.S. Investment: A Case Study of Eight Manufacturing Industries," 2012.

Arora, Vipin, "Aggregate Impacts of Recent U.S. Natural Gas Trends," Working Paper 43708, MPRA 2013.

_ **and Jozef Lieskovsky**, "Natural Gas and U.S. Economic Activity," *The Energy Journal*, forthcoming.

Brown, Stephen P.A. and Mine K. Yucel, "What Drives Natural Gas Prices?," *The Energy Journal*, 2008, *29* (2), 45–60.

CitiGPS, "Energy 2020: North America the New Middle East?," 2012.

Cleveland, William S., "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association*, 1979, *74* (368), 829–836.

Costello, Dave, Frederick L. Joutz, and Poonpat Leesombatpiboon, "Natural Gas Prices and Industrial Sector Responses: An Experimental Module for STIFS," 2006.

CRA, "US Manufacturing and LNG Exports: Economic Contributions to the US Economy and Impacts on US Natural Gas Prices," 2013.

Cunado, Juncal and Fernando Perez de Gracia, "Do Oil Price Shocks Matter? Evidence for Some European Countries," *Energy Economics*, 2003, *25*, 137–154.

Hamilton, James D., "This is What Happended to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics*, 1996, *38*, 215–220.

Hooker, M.A., "What Happened to the Oil Price-Macroeconomy Relationship?," *Journal of Monetary Economics*, 1996, *38*, 195–213.

ICF, "Oil and Gas Technology Innovation: How it is Spurring U.S. Economic Growth," 2012.

IHS, "The Economic and Employment Contributions of Shale Gas in the United States," 2011.

_ , "The Economic and Employment Contributions of Unconventional Gas Development in State Economies," 2012.

_ , "America's New Energy Future: Volumes 1-3," 2013.

Joutz, Frederick L. and Jose A. Villar, "The Relationship Between Crude Oil and Natural Gas Prices," Technical Paper, Energy Information Administration, Office of Oil and Gas 2006.

Kinnaman, Thomas C., "The Economic Impact of Shale Gas Extraction: A Review of Existing Studies," Other Faculty Research and Publications Paper 5, Bucknell University 2010.

Kliesen, Kevin L., "Rising Natural Gas Prices and Real Economic Activity," *Federal Reserve Bank of St. Louis Review*, 2006, *88* (6), 511–526.

Leone, Roberta, "Impact of Higher Energy Prices on the Northeast Regional Economy," *Contemporary Economic Policy*, 1982, *1* (1), 1–8.

McKinsey Global Institute, "Game Changers: Five Opportunities for US Growth and Renewal," 2013.

NERA, "Macroeconomic Impacts of LNG Exports from the United States," 2012.

Peterson Institute, "Fueling Up: The Economic Implications of the American Oil and Gas Boom," 2014.

PWC, "Shale Gas: A Renaissance in U.S. Manufacturing?," 2011.

_ , "Economic Impacts of the Oil and Natural Gas Industry on the U.S. Economy in 2011," 2012.

_ , "Shale Gas: Reshaping the U.S. Chemicals Industry," 2012.

_ , "Shale Oil: The Next Energy Revolution," 2013.

Weber, Jeremy G., "The Effects of a Natural Gas Boom on Employment and Income in Colorado, Texas, and Wyoming," *Energy Economics*, 2012, *34*, 1580–1588.

Appendix

	Mean	Standard Deviation	Skewness	Kurtosis
$\%\Delta_{RS}$	0.007	0.092	-0.943	5.32
$\%\Delta_{AC}$	-0.004	0.081	1.32	12.39
$\%\Delta_{CM}$	0.004	0.093	-2.06	7.50
$\%\Delta_{AL}$	0.006	0.010	0.166	4.83

(a) Industrial Production Indices

	Mean	Standard Deviation	Skewness	Kurtosis
$\%\Delta_{HH}$	-0.036	0.311	1.14	5.96
$\%\Delta_{IN}$	-0.062	0.189	0.313	2.79
$\%\Delta_{PP}$	0.019	0.210	0.155	3.04

(b) Natural Gas Prices

	$\%\Delta_{HH}$		$\%\Delta_{IN}$		$\%\Delta_{PP}$	
Industry	3 Lags	6 Lags	3 Lags	6 Lags	3 Lags	6 Lags
$\%\Delta_{RS}$	-0.092	-0.243***	-0.203**	-0.331***	-0.227***	-0.314***
$\%\Delta_{AC}$	-0.018	0.017	0.005	0.139	0.004	0.054
$\%\Delta_{CM}$	0.197***	0.158**	0.248**	0.147*	0.056	0.054
$\%\Delta_{AL}$	-0.120	-0.181***	-0.230***	-0.363***	-0.227***	-0.261***

(c) Spearman Rank-Order Correlations

Table 1: Descriptive statistics for the annual percentage changes in select industrial production and natural gas price series. The industrial production indices (resins [RS], agricultural chemicals [AC], cement [CM], and aluminum [AL]) and PPI prices (PP) range from 1991M01-2013M04, the Henry Hub price series (HH) is from 1994M11-2013M04, and the industrial price series (IN) is from 2002M01-2013M04. In panel (c) two asterisks indicate rejection of the null of no correlation at the 5% level and three indicates rejection at the 1% level.

	Total GC Tests	Reject No GC at 5%	Of Rejections, Not Jointly Determined
$\%\Delta_{RS}$	144	55 (38.2%)	36 (25.0%)
$\%\Delta_{AC}$	144	19 (13.2%)	15 (10.4%)
$\%\Delta_{CM}$	144	15 (10.4%)	11 (7.6%)
$\%\Delta_{AL}$	144	18 (12.5%)	18 (12.5%)

Table 2: Results from 4 different specifications of Granger causality tests for each industry and price at various lags.

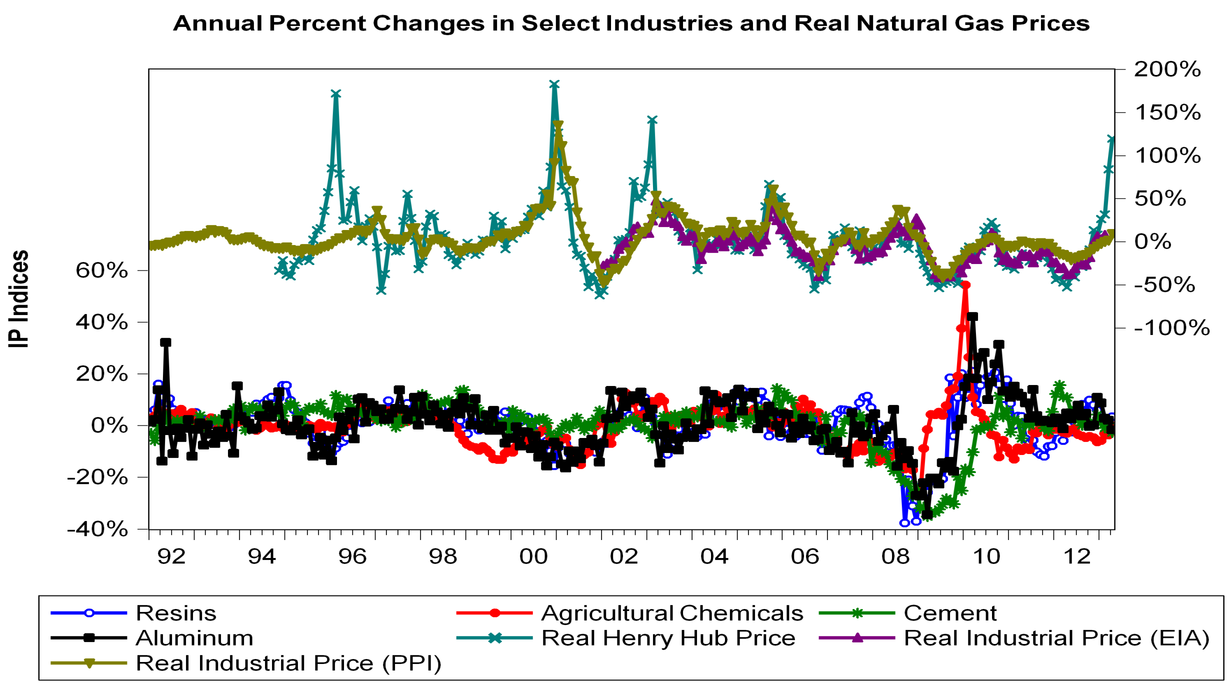

Figure 1: Annual percent changes in select series, 1991M11-2013M04.

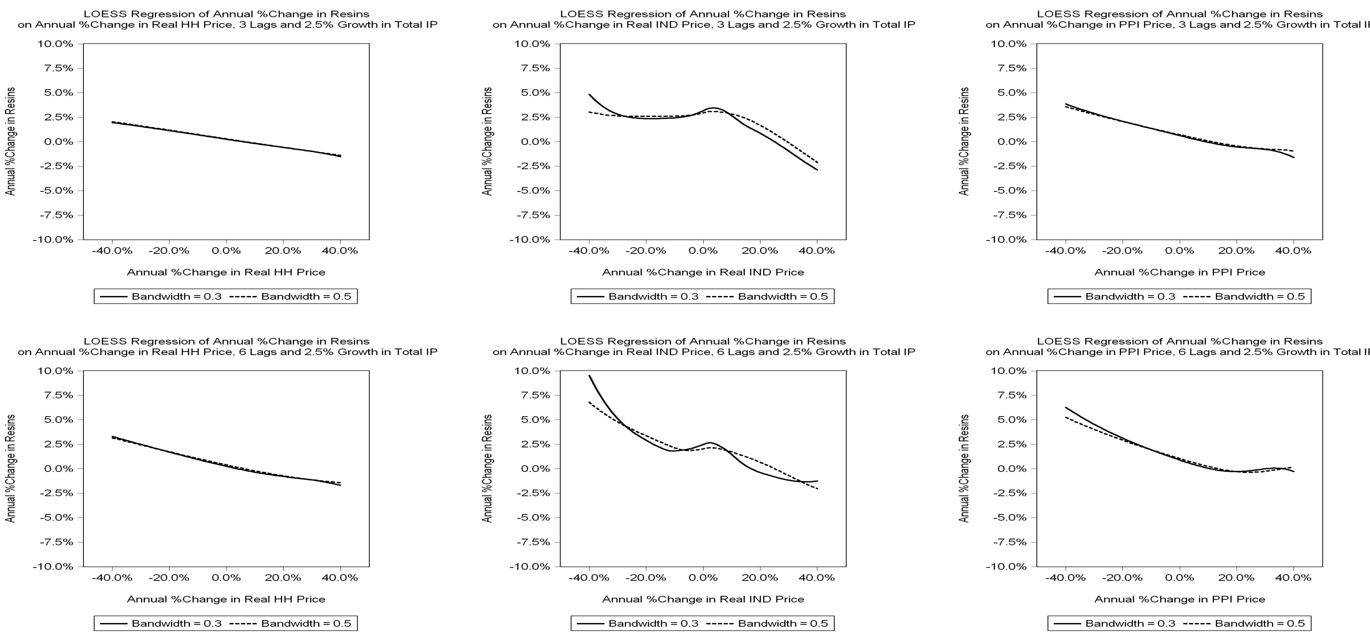

Figure 2: Non-parametric regressions of resins on each NG price and total industrial production at different lags. Regressions with HH and PPI prices range from 1994M11-2013M04, IND prices are 2002M01-2013M04.

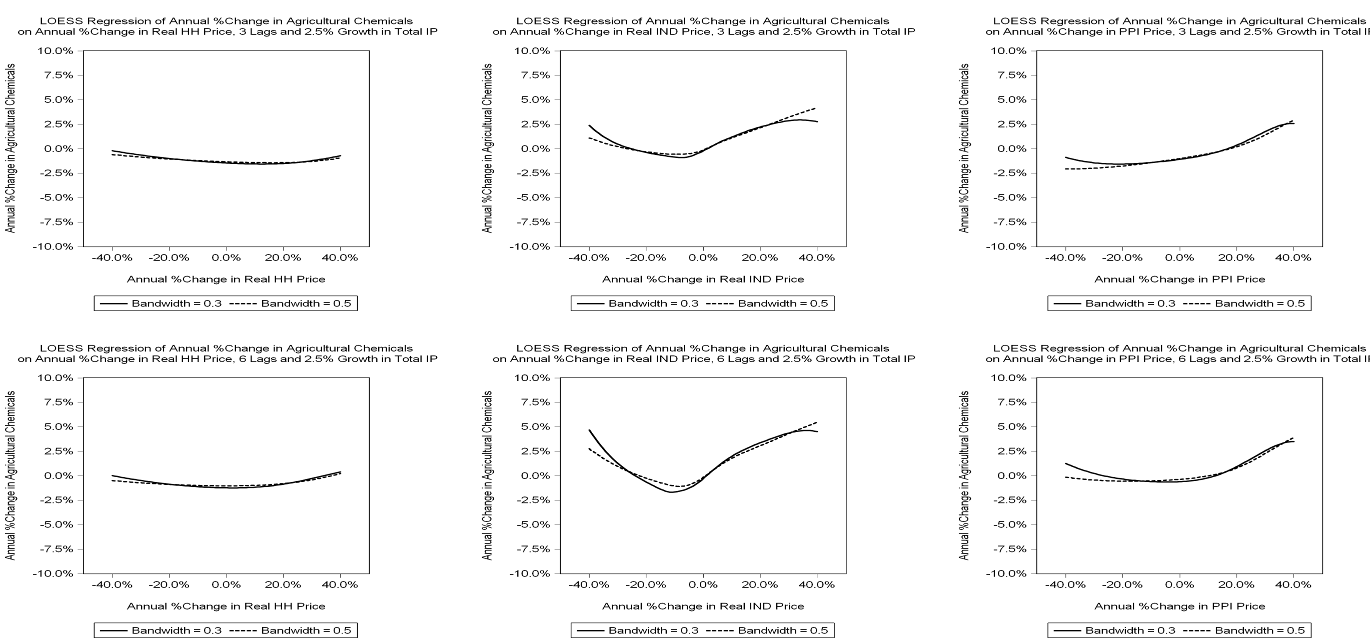

Figure 3: Non-parametric regressions of agricultural chemicals on each NG price and total industrial production at different lags. Regressions with HH and PPI prices range from 1994M11-2013M04, IND prices are 2002M01-2013M04.

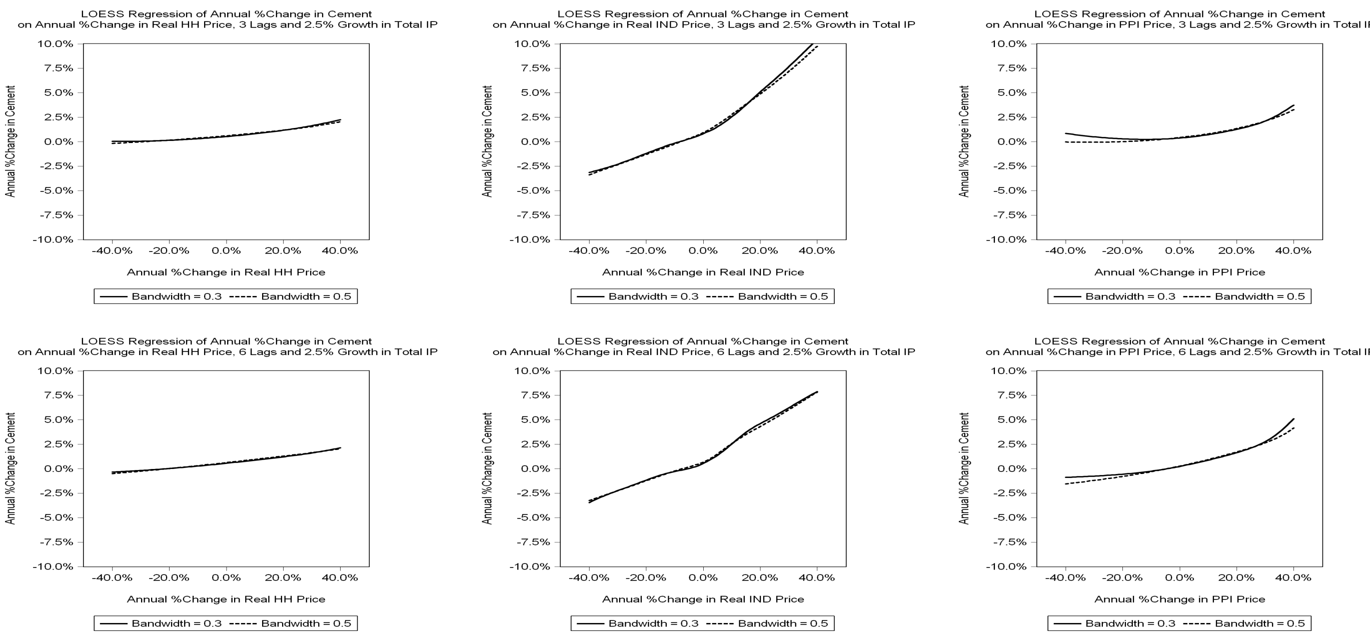

Figure 4: Non-parametric regressions of cement on each NG price and total industrial production at different lags. Regressions with HH and PPI prices range from 1994M11-2013M04, IND prices are 2002M01-2013M04.

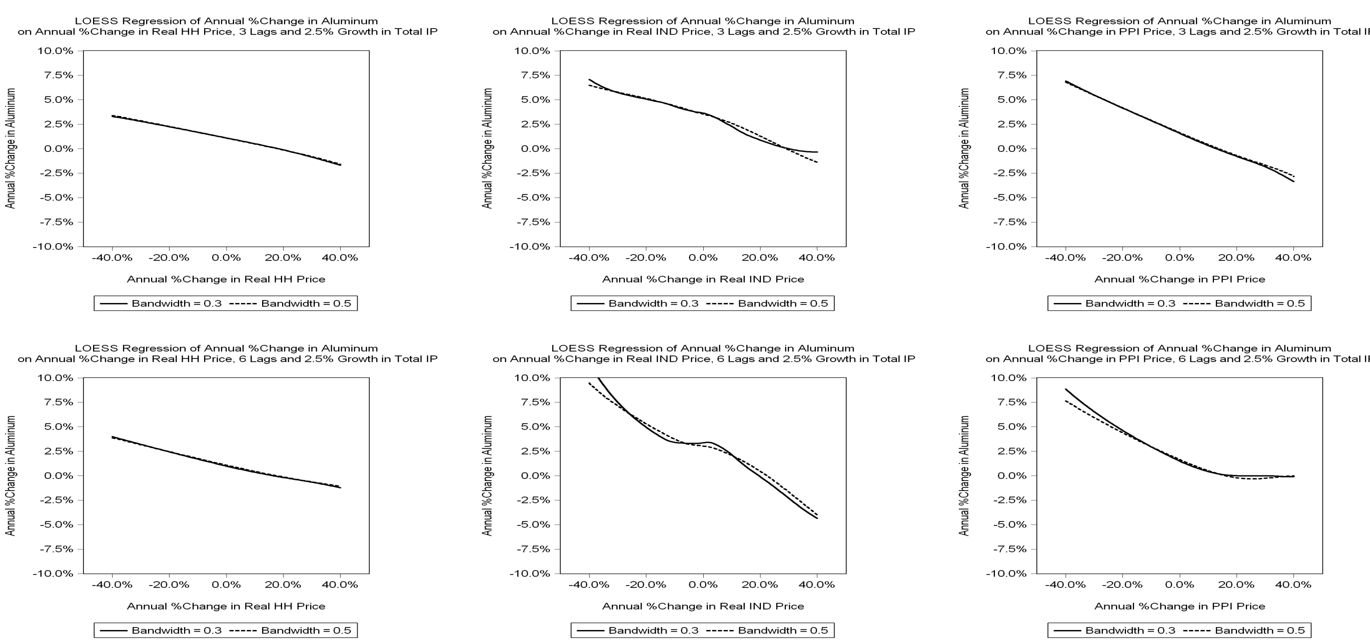

Figure 5: Non-parametric regressions of aluminum on each NG price and total industrial production at different lags. Regressions with HH and PPI prices range from 1994M11-2013M04, IND prices are 2002M01-2013M04.