Natural Gas and U.S. Economic Activity

Arora, Vipin and Lieskovsky, Jozef

September 2013

NATURAL GAS AND U.S. ECONOMIC ACTIVITY*†

Vipin Arora
U.S. Energy Information Administration

Jozef Lieskovsky
U.S. Energy Information Administration

Abstract

Previous empirical work has shown that real natural gas prices have a negligible impact on total U.S. industrial production and most of its sub-indices. We reassess these conclusions using a multivariate framework and a time-frame that includes recent developments in the U.S. natural gas market. Our results show that natural gas does affect U.S. economic activity, primarily through changes in its production. The shale gas revolution has changed this relationship - a one percentage point increase in natural gas supply raises total U.S. industrial production by more after 2008 than before.

JEL Classification: E37, F47, Q43.
Keywords: Natural gas, VAR, shale, endogenous, industrial production

1 Introduction

Relative to crude oil, comparatively little is known about the impact of the natural gas market on U.S. economic performance. This has become an important issue given recent dynamics in natural gas prices and production. Optimists continually tout the realized and potential economic benefits of the "shale-gas revolution", while many others remain unconvinced about its importance or magnitude. Implicit in either view is an assumption about the past and future macroeconomic impacts of the U.S. natural gas market, but empirical work on which to base this assumption is relatively sparse.

In this paper we evaluate the importance of the natural gas industry on U.S. macroeconomic performance. Our primary goal is to provide an empirical basis on which to judge the possible

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

†We have benefitted from the comments and suggestions of Alan Beamon, Stephen Brown, John Conti, David Daniels, Fred Joutz, Thomas Lee, Stephanie Lundquist-Arora, James Preciado, Elizabeth Sendich, Kay Smith, and Russell Tarver.

impacts of recent developments. We further highlight the channels through which the natural gas market interacts with the U.S. economy, illustrate how this relationship may be changing over time, and consider possible explanations.

While there are many studies that use macroeconomic models to evaluate the economic importance of natural gas, recent empirical work on the macroeconomic impacts of the U.S. natural gas market is limited.[1] On the empirical side, Kliesen (2006) finds that natural gas prices have historically been unable to predict total U.S. industrial production. This paper uses a single-equation framework and was completed before the large increases in natural gas production due to shale gas.[2] Costello et al. (2006) analyze the interaction between industrial natural gas prices, natural gas consumption, and industrial sector activity and conclude that industrial sector firms respond to relative energy prices, including natural gas prices. However, this paper does not focus on the aggregate impacts of the natural gas market. Similarly, Apergis and Payne (2010) and Sari et al. (2008) find a long-run relationship between natural gas consumption and economic growth in the U.S., but do not consider how the natural gas market can influence economic activity. On the other hand, Weber (2012) considers the economic impacts of the natural gas market after the shale gas boom, but on a regional level.[3]

The lack of recent empirical work on the relationship between the natural gas market and the U.S. economy is surprising given that natural gas accounts for nearly 25 percent of U.S. energy consumption, making it the second-largest source of energy behind petroleum.[4] Furthermore, consumer expenditures on natural gas have averaged about one percent of personal consumption expenditures (PCE) since 1987, and natural gas is used proportionately between the industrial, commercial and residential, and electric power sectors.[5]

Given these broad uses for natural gas, there are various ways in which the natural gas market can impact U.S. economic activity. Generally, they can be summarized as working initially through either the supply or demand sides of the economy. The most straightforward supply impact is that changes in the production of natural gas vary output in the oil and gas extraction sector, as well as associated industries. This direct change in production and its ripple through the oil and gas supply chain have been highlighted in many of the recent model-based studies on the economic

[1]Examples of studies focused on the impacts of shale gas include ACC (2011), ACC (2012), CitiGPS (2012), ICF (2012), IHS (2011), IHS (2012), PWC (2011), PWC (2012b), PWC (2012a), PWC (2013), and Kinnaman (2010) and citations therein.

[2]There is, however, a substantial amount of empirical research that explores the macroeconomic impacts of the oil market, particularly in the U.S. [see for example Hamilton (2008) and references therein]. Kilian (2009) uses a similar method to ours in quantifying the importance of different supply and demand shocks in the oil market.

[3]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.

[4]Statistics on U.S. energy consumption are from the 2011 Annual Energy Review (AER) of the U.S. Energy Information Administration (EIA). See Appendix 1 for more information.

[5]Statistics on consumer expenditures on energy are from the national accounts of U.S. Bureau of Economic Analysis (BEA). See Appendix 1 for more information.

effects of shale gas (see references in footnote 1 above).

Natural gas also influences economic activity on the supply-side of the economy through the investment of firms. The application of hydraulic fracturing and horizontal drilling has made a very large resource base available. This potential supply has led to substantial investment in the oil and gas extraction and mining support sectors, as well as other related industries. The expectation that this resource base can support lower natural gas prices for a sustained period is also leading to investment by firms outside of the oil and gas industry which rely on natural gas as an input.

Lower natural gas prices, irrespective of their cause, lower input costs for firms. These can be passed on to consumers by allowing firms to supply the same amount of goods and services at lower prices. Firms may also realize higher profits, which can lead to additional hiring, capital investment, higher dividends, or saving. Each has a follow-on impact on the economy-wide demand for goods and services. Lower prices also directly influence demand through consumers. They can raise disposable income, lower precautionary savings (or raise it in the case of a price rise), or cause consumers to change their plans for the purchases of durable goods (Kilian, 2008). Each of these result in alterations to the economy-wide demand for other goods and services stemming from the initial variation in the price of natural gas.

With these various channels in mind, we estimate a structural vector autoregression (VAR) to assess the impact of the natural gas market on U.S. economic activity. Our monthly four-variable model characterizes the supply, demand, and price of U.S. natural gas. The results are presented through impulse response analysis and variance decompositions of the model's forecast error. We also conduct sensitivity analysis on our results, including consideration of how recent developments in the U.S. natural gas market may be affecting the VAR.

The VAR model and related sensitivity analysis lead to two primary conclusions. The first is that natural gas supply changes are the primary means through which the U.S. natural gas market impacts domestic economic activity. Variations in natural gas demand for heating and power or other factors, while important for the natural gas price, do not impact economic activity in a substantial way. Our second conclusion is that the shale gas revolution has in fact changed the relationship between natural gas supply and U.S. economic activity. The responses of industrial production to the same increase in natural gas supply are larger after 2008 than before, although the size of this change remains unclear at this point.

2 VAR Model

A general VAR process can be encapsulated by a mean-zero moving average representation, without any deterministic terms (Lutkepohl, 2007):

$$\mathbf{y}_t = \sum_{j=0}^{\infty} \mathbf{B}_j \mathbf{u}_{t-j} \tag{1}$$

where $\mathbf{y}_t$ is an $N \times 1$ vector of variables, the $\mathbf{B}_j$ are $N \times N$ matrices of coefficients, and the reduced form errors ($\mathbf{u}_t$) are $N \times 1$ white noise processes with $E(\mathbf{u}_t, \mathbf{u}_t^{'}) = \mathbf{S}_u$. The coefficient matrices ($\mathbf{B}_j$) summarize the responses of the variables to the respective errors. Because $\mathbf{S}_u$ is not necessarily diagonal, the errors may be correlated across equations in the same time period. As is well-known, this can make interpretation of any responses misleading, because co-movement with other variables is not taken into account.

An equivalent representation of the moving average process with orthogonal innovations can circumvent this issue. In this case the transformed innovations will be uncorrelated by construction, so that the variance-covariance matrix of the shocks is diagonal. The identity matrix is often chosen in this case, which amounts to finding an $N \times N$ matrix $\mathbf{G}$ such that:

$$\mathbf{G}\mathbf{S}_u\mathbf{G}^{'} = \mathbf{I} \tag{2}$$

where $\mathbf{I}$ is the $N \times N$ identity matrix. The orthogonal innovations are $\boldsymbol{\epsilon}_t = \mathbf{G}\mathbf{u}_t$, so that $E(\boldsymbol{\epsilon}_t, \boldsymbol{\epsilon}_t^{\prime}) = \mathbf{G}E(\mathbf{u}_t, \mathbf{u}_t^{\prime})\mathbf{G}^{'} = \mathbf{I}$. These innovations are uncorrelated across both time and equations. In this case equation (1) can be rewritten:

$$\mathbf{y}_t = \sum_{j=0}^{\infty} \mathbf{B}_j \mathbf{G} \boldsymbol{\epsilon}_{t-j} \tag{3}$$

In this equation the $\mathbf{B}_j\mathbf{G}$ summarize the impulse responses which are plotted below. Also used in subsequent analysis to summarize model results are forecast error variance decompositions.

The forecast error variance decomposition can be reconstructed by recognizing that $\sum_{j=i}^{\infty} \mathbf{B}_j \mathbf{G} \boldsymbol{\epsilon}_{t+i-j} = E_t \mathbf{y}_{t+i}$, so that the error of the i-step ahead forecast is (Enders, 2010):

$$\mathbf{y}_{t+i} - E_t \mathbf{y}_{t+i} = \sum_{j=0}^{i-1} \mathbf{B}_j \mathbf{G} \boldsymbol{\epsilon}_{t+i-j} \tag{4}$$

From this equation we can extract the total variance in the i-step ahead forecast error of variable j, as well as the variance in the error of variable j due to variable k. The variance decompositions are reported as the fraction of the error variance in j due to k, so that the sum over all k is one.

2.1 Estimation, Data, and Identification

Our model is estimated on the annual log difference of series which range from 1993M11-2012M12 and encompass the supply, demand, and price of U.S. natural gas.[6] We use an annual difference to remove both the seasonality and trend in each variable. Unless otherwise specified, each estimation uses a lag of four months, which is chosen based on the Akaike information criterion.

Our supply variable is the change in marketed U.S. natural gas production (Δngs). We use marketed instead of gross production to exclude any gas which is used in extraction or in processing operations. And we do not differentiate between "wet" and "dry" natural gas because our interest is in the economics of production, irrespective of the type of gas supplied.

We separate the demand for natural gas between two variables. The first, total U.S. industrial production (Δipd), encompasses demand for use in the production of goods and services. This is the type of demand associated with U.S. economic activity, such as feedstock demand from a chemical firm. The second demand variable, residential natural gas demand (Δend), represents natural gas energy demand. Changes in this variable reflect variations in the demand for natural gas for heating and power purposes (possibly due to changes in the weather).

We prefer residential natural gas demand as a proxy for energy demand to the alternatives because it does not reflect changes due to economic activity. Other possible variables such as the sum of non-industrial natural gas demand do reflect changes in economic activity. For example, commercial demand for natural gas may rise because of colder weather, but it may also be higher because of overtime that requires additional power. Using only residential natural gas demand avoids this problem.

The final variable in the model is the real Henry Hub price of natural gas (Δrpg). The Henry Hub price is used because it is a marker for other natural gas prices, and changes in its value are more likely to be reflected in the prices paid by both consumers and firms than other available natural gas prices. Given these variables and consistent with the notation above, we decompose the errors and identify the shocks in the model as:

$$\mathbf{u}_t \equiv \begin{pmatrix} u_t^{\Delta ngs} \\ u_t^{\Delta ipd} \\ u_t^{\Delta end} \\ u_t^{\Delta rpg} \end{pmatrix} = \begin{bmatrix} g_{11} & 0 & 0 & 0 \\ g_{21} & g_{22} & 0 & 0 \\ g_{31} & g_{32} & g_{33} & 0 \\ g_{41} & g_{42} & g_{43} & g_{44} \end{bmatrix} \begin{pmatrix} \epsilon^{\text{“Supply” shock}} \\ \epsilon^{\text{“Economic Activity” shock}} \\ \epsilon^{\text{“Energy Demand” shock}} \\ \epsilon^{\text{“Other” shock}} \end{pmatrix} \tag{5}$$

For clarity we use quotations around the description of each shock. The first shock, a "Supply"

[6] See Appendix 1 for full details on the data.

shock, is an unexpected change in U.S. marketed natural gas supply. One concern with interpreting this as a “Supply” shock is that the supply of natural gas from storage is omitted. Unexpected changes in the supply of natural gas from storage are reflected in the “Other” shock. An example is the supply disruption in the aftermath of Hurricane Katrina. We assume that such unexpected changes can impact the remaining variables during the current month. This ordering also implies that Δngs does not respond to any of the other shocks in the current month, which reflects the costs and difficulties of changing natural gas production quickly.

The second shock, an “Economic Activity” shock, is an unexpected change in the demand for natural gas due to changes in economic activity. The sudden drop in industrial production associated with the recent U.S. financial crisis is an example. This shock can impact Δend and Δrpg in the current month. Economic activity can alter energy demand as firms alter their production plans, and natural gas prices are sensitive to expectations about future demand as well.

The third shock, an “Energy Demand” shock, is an unexpected movement in the demand for natural gas as an energy source. Such events can occur when weather is colder than expected, leading to greater demand for heating. This shock impacts only Δrpg in the current period. The final shock, an “Other” shock, represents the impact of other demand or non-demand factors on the real natural gas price. This includes the demand and supply of natural gas for storage, whether for speculative or precautionary purposes. It also includes the impact of movements in the oil price on natural gas prices.[7] These “Other” shocks cannot change either natural gas supply or demand in the current month.

3 Results

In this section we use our VAR model to understand and quantify the impact of natural gas on U.S. economic activity. We begin with the impulse responses of both the real Henry Hub price and industrial production to each of the shocks. Initially these are the responses to a one standard deviation increase in each shock, which are shown over 12 months (solid black line) along with the associated standard error bands (+/- two standard errors, dashed red lines). A one standard deviation increase is used when considering the impact of different shocks on either variable because units differ between shocks, so that raising the shocks by some percentage point leads to increases of different magnitudes. For example, an unexpected one percentage point increase in industrial production is a much larger rise than a one percentage point increase in the real Henry Hub price. Using one standard deviation helps to normalize the shocks when considering their impact on a particular variable. Percentage point increases are used when comparing the impact of one shock across different variables.

[7] There is evidence that these two variables have a time-varying relationship, and when they do move together the natural gas price tends to follow the oil price (Ramberg and Parsons, 2012).

The responses are also cumulative over the 12 months so that we are able to interpret them as the percent change in the underlying level of each variable, even though the estimated VAR is in annual log differences. We then focus on the "Supply" shocks, and show the responses of both industrial production and the real Henry Hub price to a one percentage point increase in natural gas supply.

This is followed by a discussion of the variance decompositions of industrial production and the real Henry Hub price. We then conduct sensitivity analysis on industrial production impulse responses by varying the lag length and estimating the model in levels. The next sub-section uses rolling regressions and estimations with dummy variables to consider how recent developments in U.S. natural gas markets may be affecting the VAR. The section concludes with a discussion on the plausibility of these changes.

3.1 Impulse Responses

Figure 1 shows the responses of the real Henry Hub price to unexpected increases in each of the model's four shocks. The results generally move in the direction that one would expect.

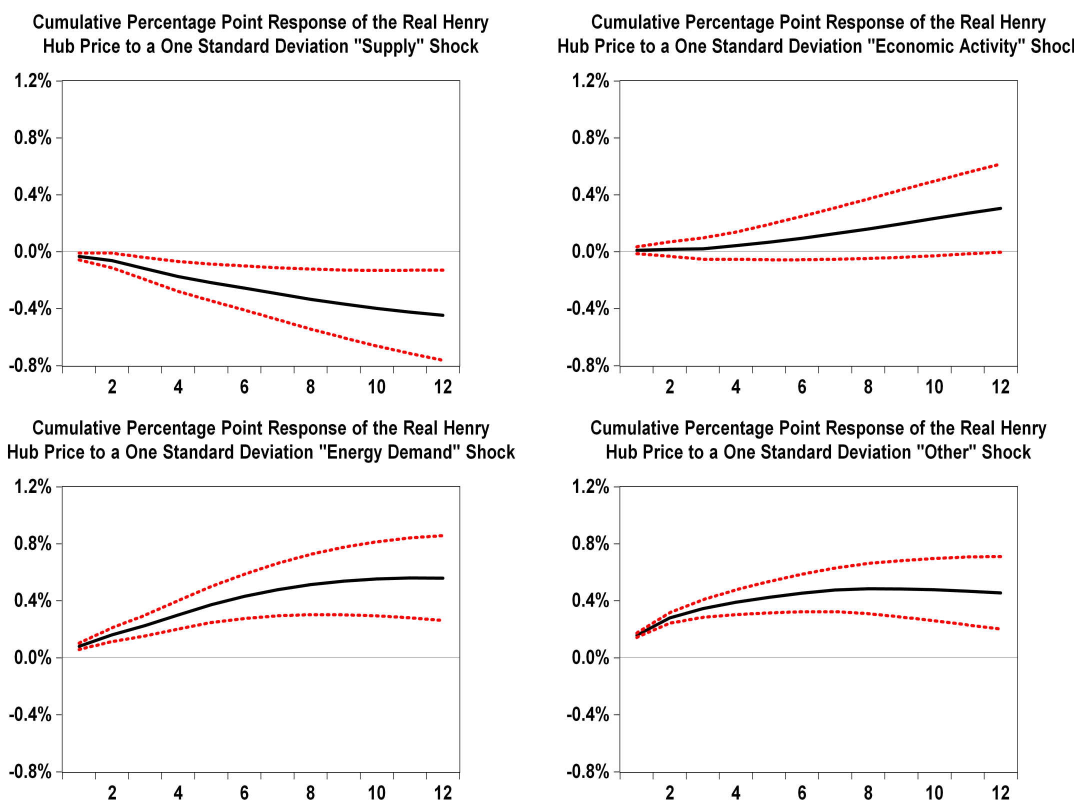

Figure 1

The top-left panel of Figure 1 shows that a rise in natural gas supply leads to a steady decrease in the real Henry Hub price over the following year. In terms of demand, the bottom-left panel of Figure 1 shows that higher demand for energy leads to an increase in the natural gas price, but the response of the price to higher economic activity is indistinguishable from zero. This may be because only a small fraction of firms use natural gas as a feedstock, or in the direct production of goods and services. But both consumers and firms use natural gas for heating and power, making changes in such demand more important for the price. An “Other” shock also leads to a higher Henry Hub price, and this may be interpreted as changes in inventories or variations in the price of oil.

Figure 2 shows the responses of industrial production to unexpected increases in each of the model’s four shocks as well. These results highlight the importance of natural gas supply for U.S. economic activity.

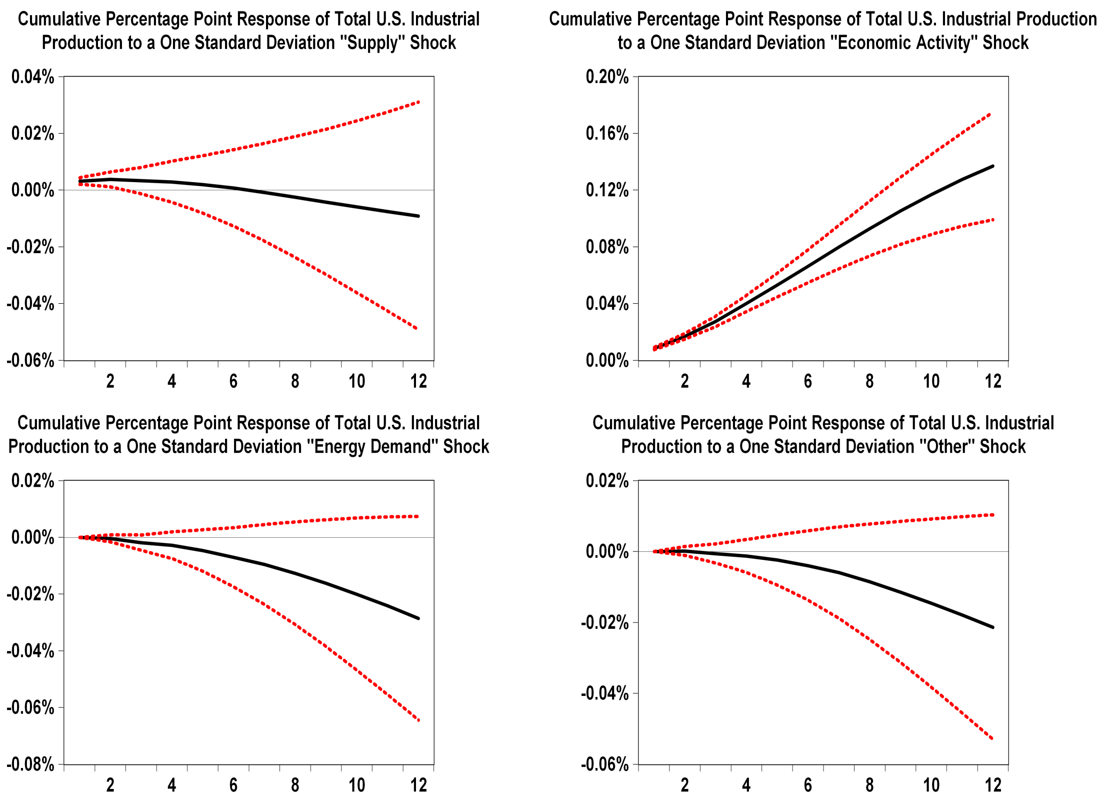

Figure 2

The top-left panel of Figure 2 shows that unexpected increases in natural gas supply lead to slightly higher industrial production. This one-to-two month increase is the result of a combination of factors. There is the direct benefit to economic activity from greater natural gas production.

This can be thought of as higher output in oil and gas extraction and mining support, as well as related activities.

As shown in the top-left panel of Figure 1, the same "Supply" shock leads to a lower natural gas price. These lower prices for consumers and firms can translate to higher disposable income, either because expenditures on natural gas fall or because goods prices are lower, and therefore increase profits for firms. Although it is unclear if these lower prices can be passed on to consumers within the one-to-two months reflected in the responses. There also has been a recent association of greater natural gas supply with investment in oil and gas extraction and mining support, which also raises economic activity.

The increase in industrial production due to its own shock is as expected. The bottom row of Figure 2 shows that the industrial production response to "Energy Demand" or "Other" shocks is small, and statistically insignificant from zero. It appears the price rises associated with either of these shocks do not translate through to reductions in industrial production. The impulse responses indicate that changes in natural gas supply are the most important factors in the natural gas market for changing industrial production.

Figure 3 illustrates this response more intuitively by showing a one percentage point increase in supply and its effect on industrial production and the real natural gas price.

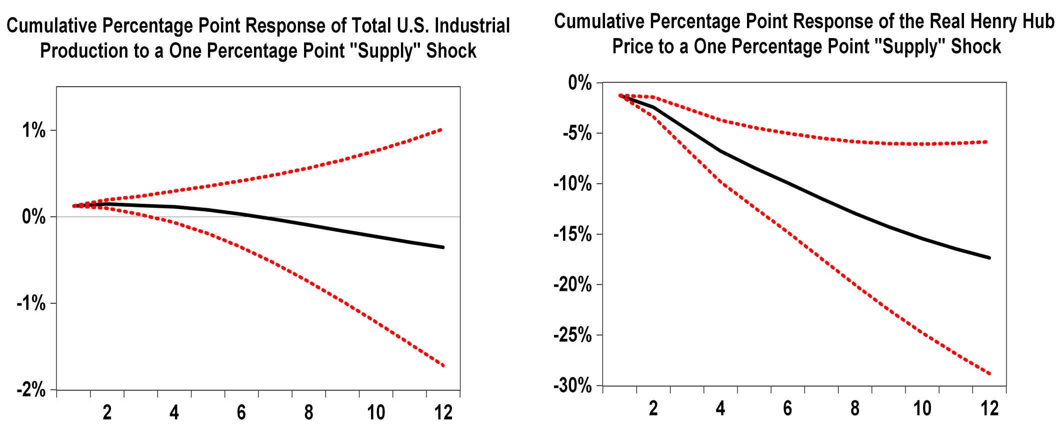

Figure 3

The shapes of each response are the same as shown in the top-left panels of Figures 1 and 2. However, the magnitude of each response is clearer in Figure 3, as is the interpretation. The average annual increase in U.S. natural gas supply has been roughly five percentage points since 2007, which corresponds to an increase of around 0.5 percentage points in industrial production and a decrease of about five percentage points in the real Henry Hub price using the results of the impulse responses.

3.2 Variance Decompositions

Another way to assess the impact of each shock on the variables in the model is by calculating the variance of the model's forecast error and then specifying the share of that variance due to each shock at different time horizons (variance decomposition). This is shown for the real Henry Hub price and industrial production at 1, 4, and 12 month intervals in Table 1.

Real Henry Hub Price	1M	4M	12M
"Supply" Shock	3.2	10.6	15.9
"Economic Activity" Shock	-	-	-
"Energy Demand" Shock	19.9	29.2	31.9
"Other" Shock	76.5	59.4	43.8

Industrial Production	1M	4M	12M
"Supply" Shock	12.8	2.5	-
"Economic Activity" Shock	87.2	96.5	90.0
"Energy Demand" Shock	0.0	-	-
"Other" Shock	0.0	-	-

Table 1: Percent of horizon step ahead forecast error variance of selected variable accounted for by the listed shocks. A dash means the fraction is not statistically different from zero.

For the real Henry Hub price, the "Energy Demand" and "Other" shocks are the most important. The "Other" shock accounts for over 76 percent of the variance of the forecast error for the real Henry Hub price at one month, and still over 43 percent at 12 months. The importance of the "Energy Demand" shock for the variance of the forecast error of the price grows over time, as does that of the "Supply" shock. The "Economic Activity" shock is not statistically different from zero in accounting for the variance of the forecast error of the real Henry Hub price at any time horizon.

The variance of the forecast error of industrial production is accounted for primarily by the "Supply" and "Economic Activity" shocks. The "Supply" shock, as with the impulse responses, has a substantial impact and then quickly gets smaller. The importance of the "Economic Activity" shock for the variance of the forecast error of industrial production remains dominant throughout the 12 months. In general, the variance decompositions of the real Henry Hub price and industrial production reinforce the results of each respective impulse response.

3.3 Sensitivity to Lag Length and Differencing

The results in the previous section are based on estimation of the VAR with a lag of four months and using the annual log difference of each variable. Here we evaluate the sensitivity of the industrial production responses by varying the lag length to 12 months and also by estimating the VAR in levels.

Figure 4 shows the impulse responses of industrial production (to a one standard deviation shock) when the lag length is extended to 12 months (the results hold for an 18-month lag as well).

The results are similar to those in Figure 2. The "Supply" shock leads to a small increase that is statistically significant and quickly becomes insignificant. And the response to an "Economic

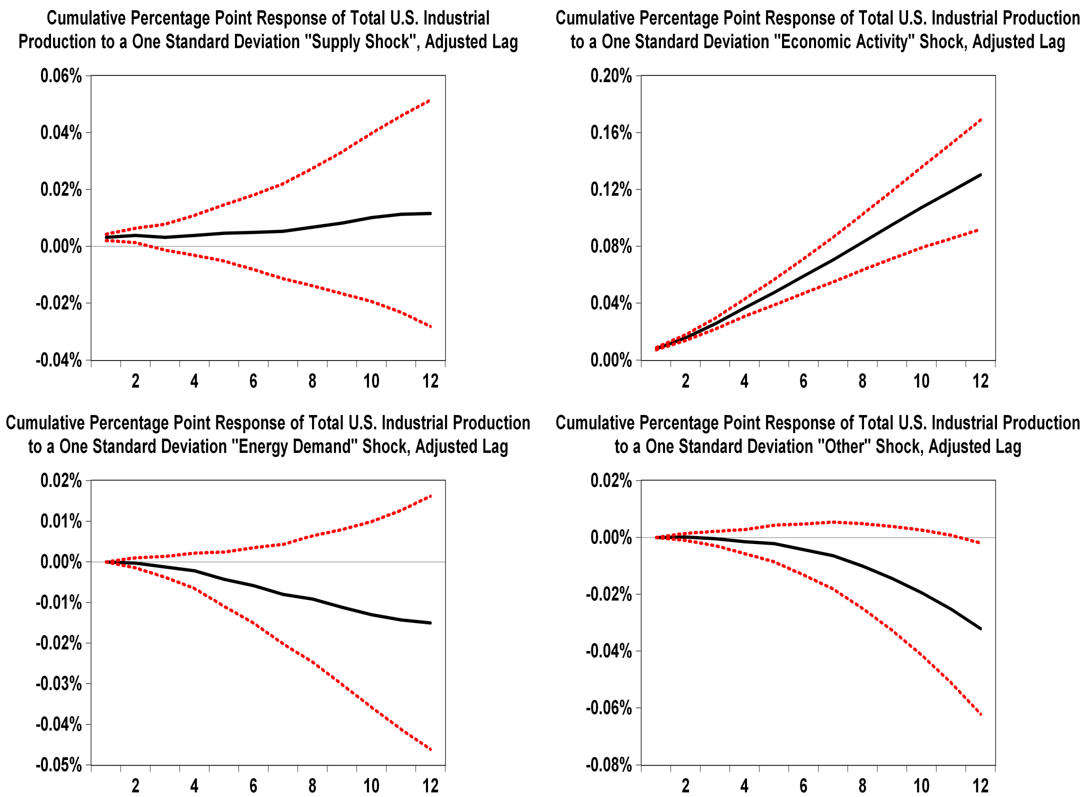

Figure 4

Activity" shock grows over the 12-month horizon. As before, the responses to "Energy Demand" or "Other" shocks are indistinguishable from zero.

Figure 5 plots the responses of industrial production using the log of each variable (with a 13-month lag, based on Akaike information criterion). Here, the cumulative responses are not necessary because we are using the levels of each variable in the estimation.

Again, the immediate response of industrial production to natural gas supply is small and quickly becomes insignificant (although the shape is different). The responses in this case are similar to the base estimation and the case with a 12-month lag. The unexpected increase in economic activity causes industrial production to rise as well, although it flattens out in this case. The responses to the remaining shocks are statistically insignificant as before.

3.4 Sensitivity to Recent Developments

The impulse responses and variance decompositions shown above assume there are no major changes, or structural breaks, in the relationship between the natural gas market and U.S. eco-

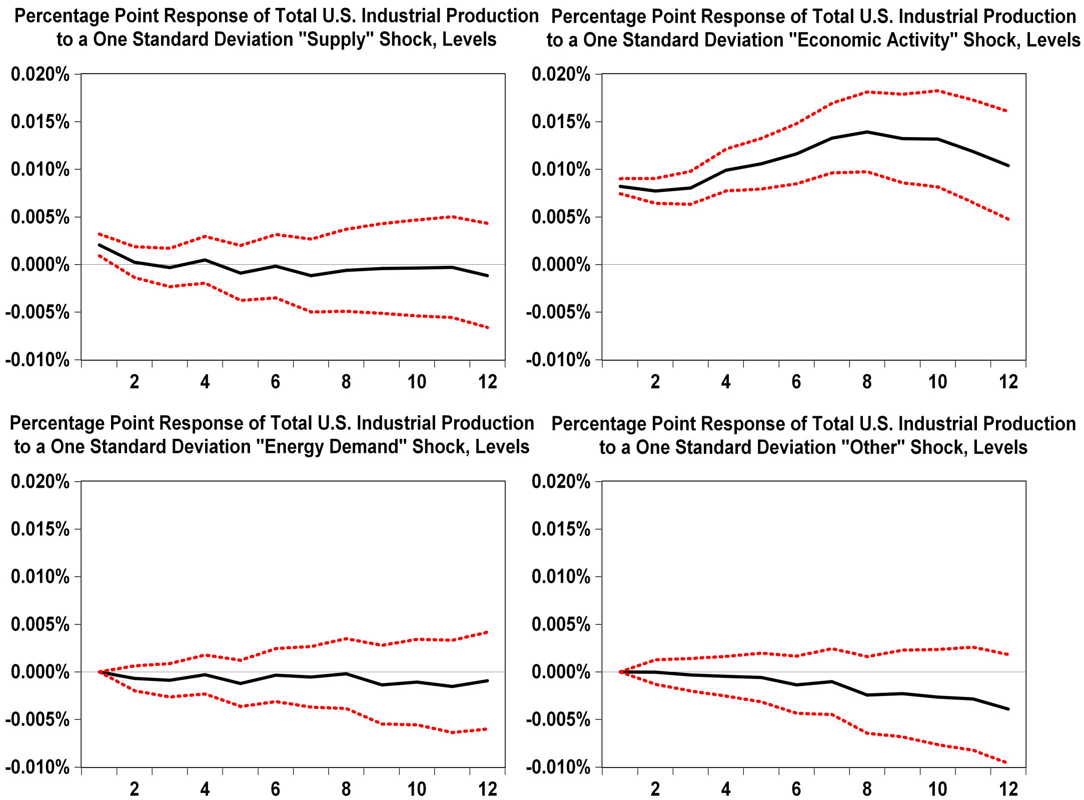

Figure 5

nomic activity throughout the sample period. This may not be true given recent developments in the U.S. natural gas market, and we consider the possibility of breaks in this sub-section. However, we do not formally test for a structural break. Our concern with tests for structural breaks is that most require an assumption that 5-15 percent of the sample, at the end, does not contain a break (Eklund et al., 2011). The recent nature of changes in the U.S. natural gas market makes this assumption difficult to verify.

Instead, we use two commonly employed techniques over a variety of break dates with a focus on the responses of industrial production. We begin by estimating the same VAR model as above using rolling regressions over different sample periods to narrow down possible break dates. The use of rolling regressions is commonly employed by forecasters in dealing with structural change (Stock and Watson, 2008). We then estimate the VAR over the full sample period with dummy variables to control for the possible break date. Our analysis considers only the industrial production responses to "Supply" shocks. The results for the remaining three shocks are indistinguishable from the responses outlined in the previous two sub-sections.[8] Based on rolling regressions, Table 2 shows the

[8]The response to an "Economic Activity" shock has a similar shape and magnitude for each rolling regression. Industrial production responses to "Energy Demand" or "Other" shocks are indistinguishable from zero in each

variance decompositions of industrial production accounted for by "Supply" shocks over different sample periods for the first three months.

Sample	1M	2M	3M
Through 2004	-	-	-
Through 2006	-	-	-
Through 2008	12.3	7.5	4.6
Through 2010	10.7	5.6	3.5
Through 2012	12.8	6.9	4.0

Table 2: Percent of horizon step ahead forecast error variance of industrial production accounted for by the natural gas "Supply" shock over the listed sample period. A dash means the fraction is not statistically different from zero.

Table 2 indicates that the "Supply" shock is unimportant for the variance of the forecast error of industrial production until the sample extends through 2008. After this time period it can account for over 10 percent of the variance of the forecast error at one month, and up to 4.6 percent after one year. These results indicate a change in the responses at sometime between the end of 2006 and the end of 2008.

To consider this further, Figure 6 shows the responses of industrial production to natural gas "Supply" shocks with the sample period ending in 2006, 2007, 2008, and 2012. Only four months are shown because each response is statistically indistinguishable from zero after this point.

There is a clear, albeit small, difference between the responses before 2008 and those after. That is, according to the model the same one percentage point increase in natural gas supply leads to initial responses of industrial production that are 0.05 percentage points higher if the sample ends in 2008 rather than 2007. Our conclusion from the variance decompositions and impulse responses based on rolling regressions is that there is likely a break in the data during 2008.

Given the results from the rolling regressions, Figure 7 shows the impulse responses of industrial production to a "Supply" shock if dummy variables are added to the VAR beginning in 2008 to control for a break.[9] The solid line is the response before 2008 and the dashed line after 2008. There is a much larger difference between the responses in this case, which indicates that the rolling regressions may be underestimating the response of industrial production to a "Supply" shock after 2008.

Using dummy variables leads to the same conclusion as with rolling regressions, that there is an amplification in the response of industrial production to unexpected changes in natural gas supply after 2008. Where the two methods diverge is in the magnitude of the amplification.

regression as well.

[9]Using F-tests we are able to reject the hypothesis that each corresponding coefficient is the same before and after the break at the 99 percent confidence level.

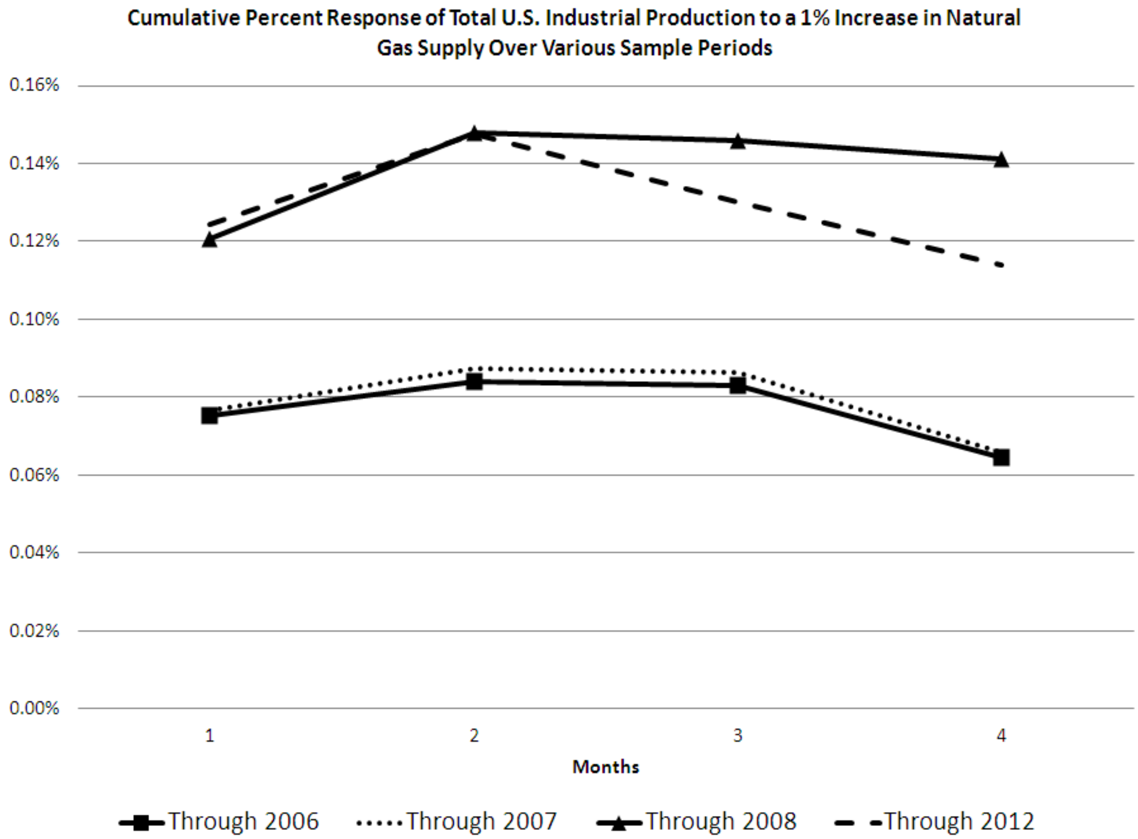

Figure 6

3.5 Are Magnified Responses Plausible?

Both the rolling regressions and dummy variable estimations indicate that the responses of industrial production to natural gas "Supply" shocks are larger after 2008 than before. In this sub-section we consider whether there is a reason that a natural gas "Supply" shock of the same size might lead to bigger industrial production responses after 2008. To this end, we analyze data on the channels through which natural gas affects U.S. economic activity. Our results are suggestive. We do not present any formal statistical analysis, mainly because the events are very recent.

The most obvious reason that industrial production responses may be larger after 2008 is that output in oil and gas extraction, mining support, and related industries is higher after this point in time. Figure 8 plots real output in oil and gas extraction and mining support over the full sample period. The vertical dashed line corresponds to 2008. Real output in these industries grows after 2008, although it is still below levels seen earlier in the sample period. This representation is also an underestimate, as it does not include additional output generated along the oil and gas supply chain. It is possible such increases help to magnify the impact of "Supply" shocks in the model, but their importance is uncertain.

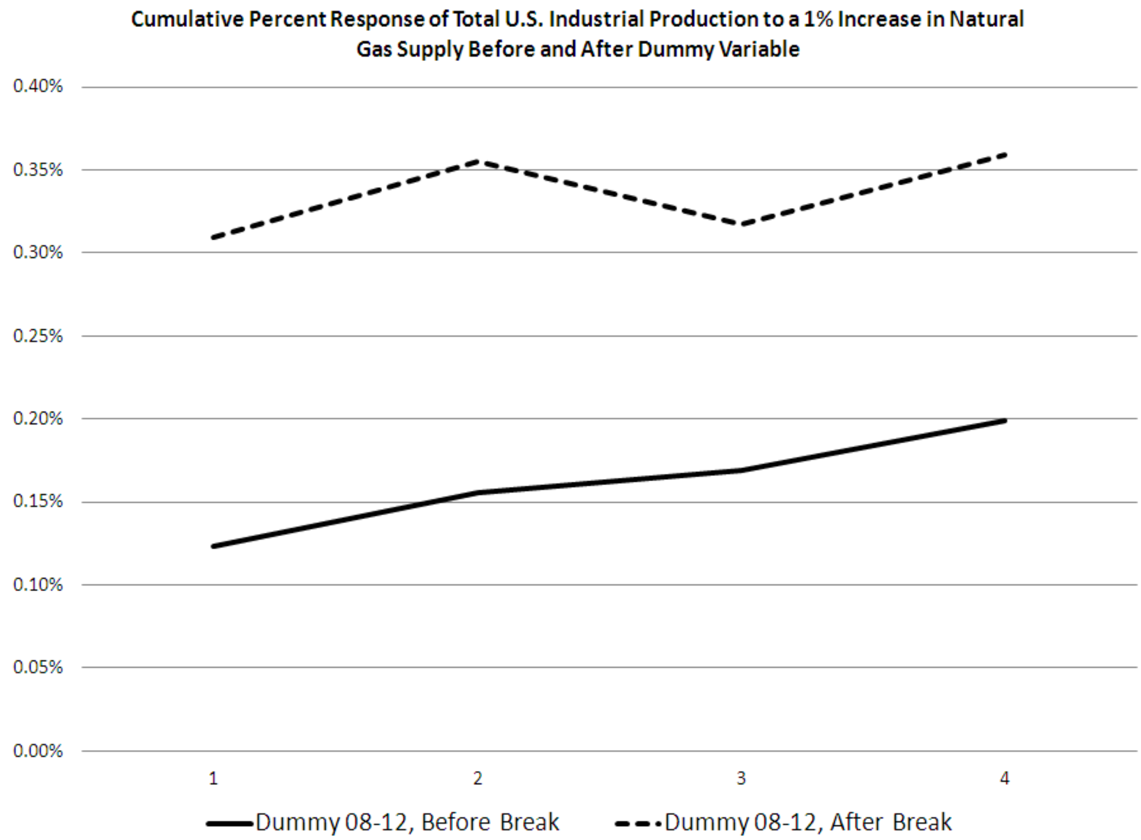

Figure 7

Lower natural gas prices, whether due to higher supply or lower demand, might also change the impact of natural gas supply on U.S. economic activity. Lower natural gas prices increase consumer disposable income, which allows for purchases of additional goods and services. This income could be higher because the direct heating and power costs paid by consumers are reduced due to falling prices. Firms might also pass on lower input costs to consumers, in which case the level of goods prices may fall. If firms do not pass on these energy savings, their profits rise, possibly leading to expansion, investment, or larger payments to owners.

Figure 9 shows the share of personal consumption expenditures (PCE) on natural gas on the left axis, and the annual percent change in the consumer price index (CPI) for commodities on the right.[10] The price level for commodities has been relatively steady before and after 2008, indicating that changes in the general price level for goods are unlikely to account for the changes in our VAR.

The share of PCE due to natural gas falls to the lowest level in our sample after 2008. However, this level is not unprecedented in the sample (1999 and 2002 are almost as low), and the share

[10]The PCE calculation is based on direct expenditures on natural gas, assuming a fifth of electricity expenditures are due to natural gas. We use a fifth because this is the share of electric power generation accounted for by natural gas.

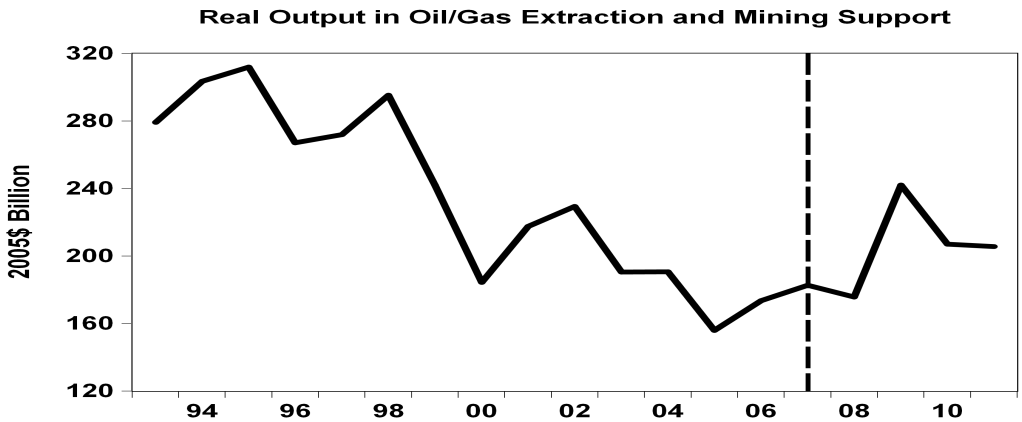

Figure 8

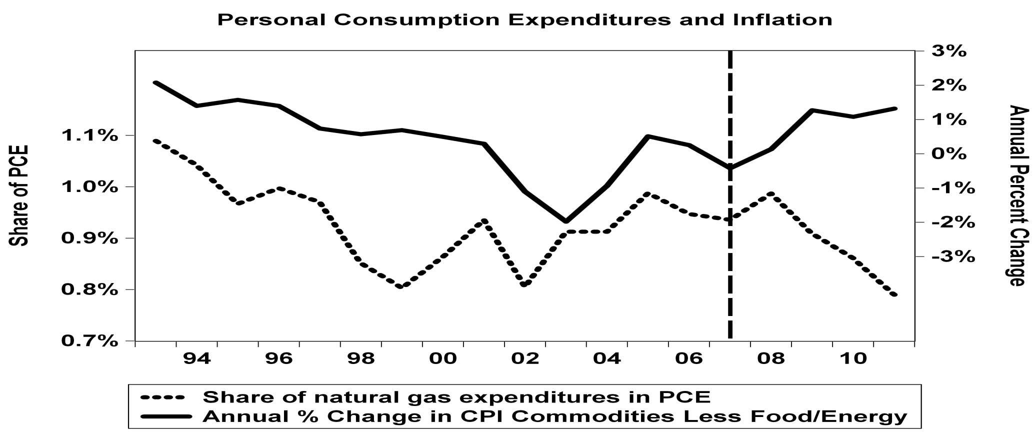

Figure 9

of PCE due to natural gas is below one percent in any case through most of the sample. The possibility that lower natural gas prices change the response of industrial production after 2008 cannot be ruled out, but its importance is unclear.

A final factor that might amplify industrial production responses to increases in natural gas supply after 2008 is the association of investment in recent years with higher levels of natural gas supply. This is because the potential size of the resource base had led firms to increase their investment in oil and gas extraction. Figure 10 shows real investment in oil and gas extraction and mining support over the sample period.

There is no clear increase in such activity after 2008 as compared to the years directly before. However, there is a clear increase beginning around 2005 and continuing through 2011. Additionally, Figure 10 shows only investment in oil and gas extraction and mining support, not investment in

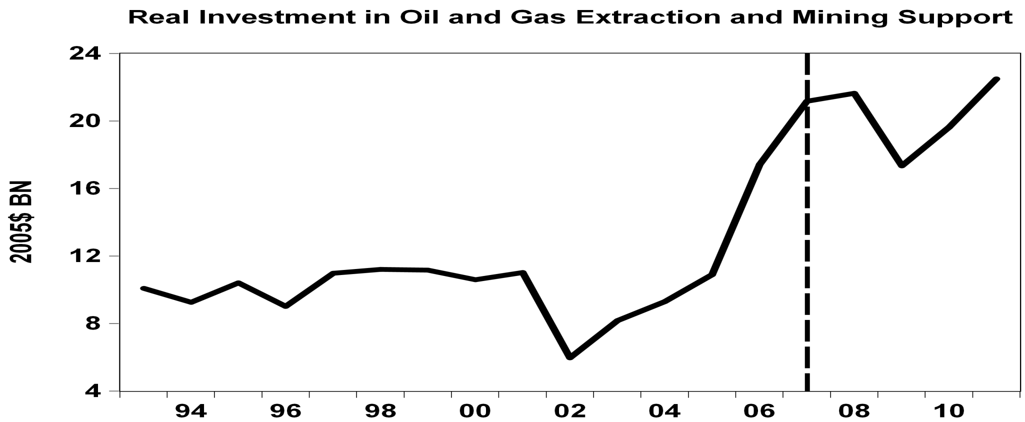

Figure 10

other associated industries along the supply chain. Nor does it include the investment by other firms, such as chemicals, which may be taking advantage of the expected lower natural gas prices for their feedstocks. Thus the actual increases in investment are likely substantially higher. As with output and prices, changes in investment might help to explain the increasing responses of industrial production to supply shocks in the VAR.

Our conclusion in light of the data is that the increasing impacts of supply shocks after 2008 obtained using the VAR model are both plausible and likely. They could be driven by a combination of greater natural gas production, lower natural gas prices, and higher investment. At this point, however, we are unable to separate out the importance of each factor.

4 Conclusion

The VAR model and related sensitivity analysis lead to two primary conclusions. The first is that natural gas supply is the primary means through which the U.S. natural gas market impacts domestic economic activity. This is supported by the quantitative results of the model and is robust to different sensitivity analysis. Our second conclusion is that the shale gas revolution has in fact changed the relationship between natural gas supply and U.S. economic activity. We show that the response of industrial production to the same size increase in natural gas supply is actually larger after 2008 than before. A look at the relevant data indicates that such a change is possible, although its magnitude is uncertain.

Both time and further research are required to know if the changes reflected in the model after 2008 are temporary or permanent. Arora (2013) uses a macroeconomic model to show that if the amplification is led by investment then it will be temporary because of the declining marginal product of capital. He argues that the natural gas industry must somehow raise economy-wide productivity to make these magnified responses more persistent. However, if the larger responses are due primarily to the increase in supply itself or to lower prices, then these responses could be permanent. This is true because it is plausible that production can continue at post-2008 levels for the foreseeable future due to the large U.S. natural gas resource base, which might result in persistently lower natural gas prices as well.

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.

Apergis, Nicholas and James E. Payne, "Natural Gas Consumption and Economic Growth: A Panel Investigation of 67 Countries," *Applied Energy*, 2010, *87* (8), 2759–2763.

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

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

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

Eklund, Jana, George Kapetanios, and Simon Price, "Forecasting in the Presence of Recent Structural Change," Working Paper 2011-23, CAMA 2011.

Enders, Walter, *Applied Econometric Time Series*, 3rd ed., Wiley, 2010.

Hamilton, James D., "Oil and the Macroeconomy," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed., Palgrage Macmillan, 2008.

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.

Kilian, Lutz, "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 2008, *46* (4), 871–909.

— , “Not All Oil Price Shocks are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 2009, *99* (3), 1053–1069.

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.

Lutkepohl, Helmut, *New Introduction to Multiple Time Series Analysis*, first ed., Berlin: Springer, 2007.

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.

Ramberg, David J. and John E. Parsons, “The Weak Tie Between Natural Gas and Oil Prices,” *The Energy Journal*, 2012, *33* (2), 13–35.

Sari, Ramazan, Bradley Ewing, and Ugur Soytas, “The Relationship Between Disaggregate Energy Consumption and Industrial Production in the United States: An ARDL Approach,” *Energy Economics*, 2008, *30*, 2302–2313.

Stock, James H. and Mark W. Watson, “Phillips Curve Inflation Forecasts,” Conference Series ; [Proceedings], Federal Reserve Bank of Boston 2008.

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 1: Data

The share of natural gas used as a fraction of total energy consumption is calculated based on the 2006 Manufacturing Energy Consumption Survey from the EIA. The data is taken from Table 1.2 of this survey, which is available at `http://www.eia.gov/emeu/mecs/mecs2006/2006tables.html`. Data on the share of personal consumption expenditures spent on select energy goods comes from the BEA. Specifically, the shares are derived from Table 2.5.5 of the National Income and Product Accounts (NIPA), which are available at http://www.bea.gov/national/index.htm#gdp at an annual frequency. Also taken from the BEA are value added in the oil and gas extraction and mining support sectors, each of which is taken from the GDP-by-Industry tables, real value added table. The investment data on these same two sectors comes from the BEA's fixed assets tables, 3.7E, and the sum of the values is deflated by the GDP deflator with the base year specified as 2005. The GDP deflator is available quarterly beginning in 1947 from the Federal Reserve Bank of St. Louis's Federal Reserve Economic Data (FRED) at http://research.stlouisfed.org/fred2/series/GDPDEF.

Data on natural gas production, the real natural gas price, and end-use consumption of natural gas are taken from the EIA. Natural gas production is marketed U.S. natural gas withdrawals in millions of cubic feet, and is available at a monthly frequency from 1980M01 at http://www.eia.gov/dnav/ng/hist/n9010us2m.htm. The real natural gas price is the Henry Hub price deflated by the U.S. producer price index (PPI) for fuels and related products and power (WP05). The Henry Hub price is available at a monthly frequency from 1993M11 and comes from http://research.stlouisfed.org/fred2/series/GASPRICE/. The PPI is available monthly from 1926M01 and is taken from the Federal Reserve Bank of St. Louis's Federal Reserve Economic Data (FRED) at http://research.stlouisfed.org/fred2/series/PPIENG. End-use consumption of natural gas is available from the EIA's monthly energy review. Data on consumption of natural gas by end-use sector is available in Table 4.1 of this document, and the monthly historical data can be found at http://www.eia.gov/totalenergy/data/monthly/#naturalgas in billions of cubic feet. In calculating the natural gas demand for use as energy, we use natural gas consumed in the residential end-use sector.

The total U.S. industrial production index is taken from the Board of Governors of the Federal Reserve, and is available at http://www.federalreserve.gov/releases/g17/download.htm at a monthly frequency from 1967 onwards. The total index includes manufacturing, mining, and utilities.