Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?

Lutz Kilian
University of Michigan and CEPR

Bruce Hicks
University of Michigan

March 14, 2010

Abstract: Recently developed structural models of the global crude oil market imply that the surge in the real price of oil between mid-2003 and mid-2008 was driven by repeated positive shocks to the demand for all industrial commodities, reflecting unexpectedly high growth mainly in emerging Asia. We evaluate this proposition using an alternative data source and a different econometric methodology. Rather than inferring demand shocks from an econometric model, we utilize a direct measure of global demand shocks based on revisions of professional real GDP growth forecasts. We show that recent forecast surprises were associated primarily with unexpected growth in emerging economies (in conjunction with much smaller positive GDP-weighted forecast surprises in the major industrialized economies), that markets were repeatedly surprised by the strength of this growth, that these surprises were associated with a hump-shaped response of the real price of oil that reaches its peak after 12 to 16 months, and that news about global growth predict much of the surge in the real price of oil from mid-2003 until mid-2008 and much of its subsequent decline.

Key words: Oil price; Global Real Activity; Demand; News; Shocks; Forecast Revisions; EIU.

JEL: C42, C53, Q43

Acknowledgements: We thank Ana-María Herrera and two anonymous referees for helpful comments.

Correspondence to: Lutz Kilian, Department of Economics, University of Michigan, 238 Lorch Hall, 611 Tappan Street, Ann Arbor, MI 48109-1220, USA. Email: lkilian@umich.edu. Fax: (734) 764-2769. Phone: (734) 647-5612.

1. Introduction

A central question for macroeconomists and financial analysts alike is what caused the dramatic surge in the real price of oil between 2003 and mid-2008. The structural VAR model of Kilian (2009a) implies that this surge was driven by repeated positive shocks to the demand for industrial commodities including crude oil.[1] This model relies on the use of a proxy for fluctuations in global real economic activity based on dry cargo ocean shipping freight rates. Similar measures of economic activity such as the Baltic Dry Cargo Index are also in use among market practitioners. Further analysis in Kilian (2009a) based on a linearly detrended index of OECD industrial production as an alternative proxy for global real economic activity suggests that the unexpected increase in the demand for oil after 2002 was not driven primarily by unexpectedly high growth in the OECD, but to a large extent by unexpected growth from countries outside of the OECD. This finding is consistent with the widespread perception that much of the recent boom in industrial commodity markets was driven by the economic transformation of countries in emerging Asia such as China and India.[2]

At first sight it may strain credulity that markets would have been repeatedly surprised by high growth in emerging Asia, as suggested by the econometric model, rather than adjusting their expectations early on when it became apparent that the emerging Asian economies were booming. In this paper, we show that this central implication of the model is consistent with independent evidence based on professional forecasts for real economic growth in China and other countries. Based on the data provided by the Economist Intelligence Unit we first document that, starting in mid-2003, forecasters were repeatedly surprised by high economic

[1] This type of a model has been employed in a variety of contexts to study the effect of oil demand and oil supply shocks on macroeconomic aggregates and financial markets (see, e.g., Alquist and Kilian 2010; Kilian 2008; Kilian 2010; Kilian and Park 2009; Kilian, Rebucci and Spatafora 2009).
[2] See, e.g., Considine (2006); Kilian (2009a,b); Hamilton (2009a,b); Smith (2009).

growth in emerging economies. In contrast, forecast surprises about OECD real economic growth were much smaller.

Second, we construct estimates of the response of the real price of crude oil to weighted surprises in professional real GDP forecasts. We exploit the fact that suitably weighted real GDP forecast surprises can be treated as news shocks for the global crude oil market.[3] We show that the response of the real price of oil to such news shocks is similar to the response to broad-based global demand shocks in industrial commodity markets in the structural VAR model of Kilian (2009a). Unexpected growth in China, for example, is associated with a large hump shaped response that builds slowly and peaks after about one year. The same regressions for an aggregate of the United States, Germany and Japan yield an increase in the real price of oil that peaks after 16 months. Third, a historical decomposition shows that growth surprises in emerging economies as well as advanced economies jointly explain much of the rise and decline of the real price of oil between 2000.12 and 2008.12, underscoring the importance of fluctuations in global real economic activity for the real price of oil. Fourth, estimates of cross-autocorrelations show that real GDP forecast news about India and China tends to lead innovations to the global demand for industrial commodities.

We conclude that unexpected growth in emerging economies after 2003 played a central role in driving up the real price of oil until mid-2008 and was reinforced by smaller positive growth surprises in major OECD economies. Likewise, much of the decline in the real price of oil is explained by large negative growth shocks since mid-2008 both in emerging and in

[3] The link between innovations to global real GDP and the real price of oil is discussed, for example, in Barsky and Kilian (2002). Our econometric methodology is based on the work of Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Vega (2003, 2007) and Faust, Rogers, Wang, and Wright (2007), for example. Kilian and Vega (2010) conduct a similar analysis using high frequency U.S. macroeconomic news measures. For a related analysis of oil inventory surprises see Arsenau, Beechey, and Vigfusson (2008). The news shocks in this literature are observable to the econometrician and should not be confused with unobservable "information" shocks, as discussed in the recent macroeconomic literature, which must be inferred from the data by the econometrician.

advanced economies. The remainder of the paper is organized as follows. Section 2 discusses the data on real GDP forecast surprises and the evolution of the real GDP weights. Section 3 presents the econometric methodology and impulse response estimates. In section 4 we discuss the relationship between professional real GDP forecast revisions and innovations to the global demand for industrial commodities identified in structural VAR models and show that the data are consistent with a number of complementary explanations of this link. Section 5 contains the concluding remarks.

2. Forecast Surprises

2.1. Data Construction

We construct measures of exogenous shocks to real activity based on forecasts of real activity provided by the Economist Intelligence Unit (EIU).[4] The EIU is one of the leading providers of such forecasts. Country-specific forecasts of annual real GDP growth for the current and future years are available every month. The sample period is 2000.11-2008.12.[5] We define revisions of forecasts of real GDP growth as a forecast surprise or news shock.[6] Let $F_{i,t}$ denote the EIU forecast of annual real GDP growth for the current year or the next year. We focus on the one-year forecast horizon because one-year forecasts are more reliable and watched more closely by market participants and because there is much less variability in forecast revisions at longer horizons. We follow the EIU in treating current annual growth as unknown for January through September of that calendar year. Starting in October of every year, we follow the EIU in focusing on forecasts for the following calendar year. A news shock then can be defined as the

[4] The data are proprietary and available from the EIU at www.eiu.com
[5] No monthly data are available prior to this date.
[6] An alternative approach would be to compare forecast to ex post realizations of ex post revised data. Since the lags with which such data are available are long and variable, since the quality of the final data for China and India is questionable, and since we are interested in measuring the news component of forecast announcements in real time, defining news in terms of forecast revisions is more natural.

forecast revision:

$$N_{i,t} = F_{i,t} - F_{i,t-1},$$

where i denotes the country and t denotes the current month. Since the news shocks already are expressed in annualized percent growth rates we do not standardize them. However, a given country's news shock is weighted by the time-varying share $\omega_{i,t}$ of the country's purchasing-power adjusted real GDP in purchasing power adjusted world real GDP:

$$S_{i,t} = N_{i,t}\omega_{i,t}$$

This approach helps us capture the growing importance of economies such as China in the world economy over our sample period. The share data are constructed from purchasing power adjusted real GDP data in the Penn World Table 6.2.[7] Since the table ends in 2003, we extrapolate the shares for 2004 through 2009 from the average growth rate of each country's share over 1996-2003.[8] Given the approximate linearity of the time path of the shares, this approach should provide a good approximation. Since the Penn World Table is annual, we linearly interpolate the shares to obtain the monthly weights used in constructing $S_{i,t}$. Similarly, in constructing aggregates of news shocks across countries we use country weights reflecting purchasing-power adjusted real GDP estimates:

$$\begin{aligned} S_{1,t} &= \sum_{i=1}^{2} N_{i,t}\omega_{i,t} \\ S_{2,t} &= \sum_{i=3}^{5} N_{i,t}\omega_{i,t} \end{aligned},$$

where $i = 1$ denotes China, $i = 2$ stands for India, $i = 3$ for the United States, $i = 4$ for Germany, and $i = 5$ for Japan.

[7] These data are available from the Center for International Comparisons at the University of Pennsylvania at pwt.econ.upenn.edu.
[8] Although the Penn World Table 6.2 includes some data for 2004, there are no real GDP data for the world as a whole.

2.2. Data Analysis

A testable implication of the analysis in Kilian (2009a) is that real GDP forecast surprises in emerging Asia should be positive on average and should be much higher than in OECD economies starting in 2003. Table 1 supports that view. We focus on the four largest emerging economies (Brazil, Russia, India and China), often collectively referred to as the BRIC countries, and the three largest OECD economies (the United States, Germany and Japan). Table 1 shows that of these economies only Russia exhibited strong unexpected growth in 2000.12-2003.5. All other economies experienced (often large) negative growth shocks. In contrast, between 2003.6 and 2008.6, the world economy accelerated unexpectedly and – with the exception of Germany – across the board. The largest growth shocks were recorded in China, Russia, and Japan. In addition, India, Brazil, and the United States experienced smaller positive growth shocks. After 2008.6, growth collapsed unexpectedly in all countries with the exception of Brazil, as did the real price of oil. The largest forecast revisions occurred in Germany, Japan and Russia (in that order), but even China, and India experienced large negative forecast surprises. This evidence is consistent with the view that the oil price shock of 2003-2008 was driven at least in part by an unexpected acceleration and deceleration of world economic activity.

Table 1 shows that unexpected growth in emerging economies played an important role in this event, but not at the complete exclusion of growth shocks in advanced economies. Japan's unexpected recovery starting in mid-2003 certainly was a factor. More generally, the relative importance of different countries evolved over time. For example, growth shocks in emerging economies such as China were dominant between mid-2003 and mid-2008, whereas the decline in the real price of oil since mid-2008 was associated as much with negative growth shocks in OECD economies as in emerging Asia.

Many of the forecast surprises during 2003.6-2008.5 were quite sizable. For example, the average monthly forecast surprise about Chinese GDP growth was +0.12 percentage points of annual growth. While a shock of this magnitude may seem modest by itself, successive shocks of this magnitude over several years have the potential to trigger large adjustments in the demand for industrial commodities in general and crude oil in particular. Economic theory suggests that unexpectedly high economic growth without a commensurate increase in global oil supplies should be associated with increases in the real price of oil, as global demand for crude oil grows (see, e.g., Barsky and Kilian 2002).

3. Estimating the Effect of Forecast Surprises on the Real Price of Crude Oil

A natural question is how much of the observed movements in the real price of oil are explained by the forecast surprises documented in section 2. Since response estimates for individual countries can be erratic and since there is reason to believe that emerging economies have a larger industrial sector than advanced economies, we focus on two broad aggregates of countries. The emerging economy aggregate includes China and India, which jointly account for 23.1% of world GDP in 2008 (see Table 2). It may have seemed natural to include the remaining BRIC countries as well. We exclude Russia from the set of emerging economies because it is unclear whether Russian growth shocks are exogenous with respect to the real price of oil, given the dependence of the Russian economy on foreign exchange earnings from oil and gas exports. We exclude Brazil because that economy is effectively decoupled from global crude oil markets to the extent that much of its energy needs are satisfied by domestic ethanol production. We note, however, that the estimated responses for the aggregate for China and India shown below would be very similar, if we included Brazil and Russia among the emerging economies. The aggregate of the advanced economies consists of the United States, Germany and Japan and accounts for

28.3% of world GDP in 2008.

While instructive, the forecast surprises shown in section 2 do not account for the increasing weight of China and India in the world economy in recent years, and the declining weight of OECD economies. Table 2 shows that the combined weight of the United States, Germany and Japan has declined from 33.6% in 1996 to 28.3% in 2009. The weight of all OECD economies has fallen from 58.4% to 51.0%. At the same time, the combined weight of China and India has risen from 14.4% to 23.1%, with China alone accounting for almost 16% of the world economy in 2009. In assessing the impact the forecast surprises had on the crude oil market, it is important that we control for the increasing weight of China and India in the world economy in recent years.

Table 3 assesses the relative importance of the forecast surprises for the world economy. Of particular interest is the column for 2003.6-2008.6, when the real price of oil surged. It shows that the weighted U.S. forecast surprises during this period were small on a global scale, as were the weighted forecast surprises for Germany and Japan combined. In contrast, the weighted forecast surprises for China and India combined were 3.5 times as large as for the United States, 3 times as large as for Germany and Japan combined and almost twice as large as for the entire OECD. This evidence is consistent with the view that unexpected growth in emerging Asia was an important contributor to the shift in demand for industrial commodities including crude oil.

Figure 1 plots the weighted forecast surprises since 2001 for China and India combined and for the aggregate of the United States, Germany and Japan. The correlation of the two forecast error series is 0.29. A possible interpretation is that stronger than expected growth in emerging Asia also stimulates industrialized economies that export machinery and other manufactured goods to emerging Asia. The magnitude of this contemporaneous correlation is not

statistically significant, however. We treat the series in Figure 1 as measures of exogenous shocks to global real economic activity, consistent with the standard approach in the news shock literature.

Data for the real price of oil are obtained by deflating the U.S. refiners' acquisition cost of imported crude oil by the U.S. consumer price index for all urban consumers.[9] We relate the percent change in the real price of oil to current and lagged news shocks:

$$\begin{aligned} \Delta rpoil_t &= \alpha_1 + \sum_{j=0}^{18} \beta_{1,j} S_{1,t-j} + \varepsilon_{1,t} \\ \Delta rpoil_t &= \alpha_2 + \sum_{j=0}^{18} \beta_{2,j} S_{2,t-j} + \varepsilon_{2,t}, \end{aligned} \tag{1}$$

where $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are possibly serially correlated. Following the literature on news regressions, the baseline regression estimates are based only on dates t on which a nonzero forecast surprise occurred. The parameter $\beta_{1,j}$, for example, measures the response of $\Delta rpoil_{t+h}$, $h = 0, 1, 2, \ldots,$ to a unit news shock for China and India at date t after j months. The parameter $\beta_{2,j}$ indicates the corresponding response to a unit news shock for the aggregate of the United States, Germany and Japan. An estimate of $\beta_{1,j} = 0.1$ specifically would imply that a forecast revision for the aggregate of China and India of one percentage point would cause an increase in the real price of oil by 0.1%. Since actual forecast surprises are much smaller, we scale all responses to represent the effects of a +0.1 percentage point shock in N_{it}. In practice, we report the cumulative responses with one standard error bands obtained by the block bootstrap method (see, e.g., Berkowitz, Birgean, and Kilian 1999).[10]

[9] The oil price data were obtained from the *Monthly Energy Review* of the Energy Information Administration at http://www.eia.doe.gov. The CPI data are from the FRED database.
[10] Results shown are for a block size of 8. Almost identical results are obtained with block sizes of 4 and 12.

The two panels of Figure 2a show the dynamic response of the price of oil to the news shocks $S_{1,t}$ and $S_{2,t}$. In the first panel, a revision of annual real GDP growth forecasts for China and India by 0.1 percentage points raises the real price of oil by about five percent. The response is hump shaped with a peak after 10 months. The price response is quite persistent and most statistically significant in months 9 through 15.[11] The same type of shock in the second panel induces a price increase of almost the same magnitude, with a peak after 16 months, but the hump shape is less pronounced and the response is barely statistically significant in months 8 through 13.[12] Figure 2b shows that similar results are obtained even if we include periods of zero forecast surprises in the regressions. The main difference is that the peak response to $S_{1,t}$ occurs after 14 rather than 10 months.

It is important to stress that there is no reason for the real price of oil to respond immediately to real GDP forecast revisions on impact. Related work by Kilian and Vega (2010) shows that there is no evidence of feedback from daily U.S. macroeconomic news to the price of oil either on impact or within the month, as one would expect if the price of oil responded instantaneously to all news. This finding is not surprising as even stock prices may adjust to demand shocks gradually over time (see, e.g., Kilian and Park 2009). Crude oil, moreover, like most mineral products differs from other goods and services in that it is both storable and relatively homogeneous. As a result, it is a hybrid of *assets* – whose price is determined by supply of and demand for *stocks* – and *goods* whose prices are determined by *flow* supply and *flow* demand matter (see, e..g., Frenkel and Rose 2009). This interpretation is also consistent

[11] Similar results hold for regressions on China shocks alone. That response is slightly larger than the response to for the aggregate of China and India, but otherwise similar.

[12] Regressions for broader forecast aggregates including all OECD economies yield a similar hump shaped response with a peak after 11 months, but the response is very imprecisely estimated. This result may reflect the lower quality of forecasts (and real GDP weights) for many smaller OECD economies as well as a preoccupation of markets with the real GDP forecasts for the major economies. We therefore concentrate on the three largest OECD economies.

with recent work on oil demand shocks in Kilian (2009a).

The upper panel of Figure 2c shows for comparison the response to a positive shock to global demand for all industrial commodities estimated from the Kilian (2009a) structural VAR model. This VAR response is the direct analogue of the response shown in the lower panel of Figure 2c, which is based on regression model (1) applied to a suitably weighted aggregate of the three OECD economies, China and India. Although by construction the magnitude of the shock is not the same, the overall pattern of the impulse response is quite similar.[13] The responses are weakly hump shaped and persistent. They build slowly with a peak after 12 or 13 months for the VAR estimate and somewhere between 10 and 16 months for the news regression estimate. These aggregate results, of course, ignore inherent differences between shocks originating in emerging economies and in OECD economies.

One immediate implication of our analysis is that for the same shock $S_{k,t}, k \in \{1,2\}$, the response to $S_{2,t}$ would be somewhat smaller in magnitude than the response to $S_{1,t}$. This result is expected. Since much of the world's industrial production has moved to emerging Asia, a given surprise about real GDP growth all else equal is associated with a larger increase in the demand for industrial commodities, if the shock emanates from that region. More generally, it is not real GDP that is driving the demand for industrial commodities, but only certain components of real GDP. For example, growth in the service sector (which is more dominant in OECD economies) is associated with less additional demand for industrial commodities than growth in manufacturing. In addition, the same growth in manufacturing may lead to a much larger surge in the demand for industrial raw materials, iron ore, cement and crude oil in countries such as

[13] Note that it does not make sense to normalize the shock to be of the same numerical magnitude, as the shocks underlying Figure 2c are measured in different units. In particular, a shock of unity for the global aggregate demand shock would not correspond to a unit innovation in real GDP in the world or in any one country or group of countries. Our main concern here is whether the shape of the response function is similar rather than the magnitude.

China that have to build the infrastructure to sustain the growth in manufacturing, compared with industrialized economies with an existing infrastructure.[14]

Having estimated the responses to a one-time real GDP forecast error, we now assess the overall importance of news about real economic activity for the real price of oil since mid-2001 based on the fitted value of the regression (1). We construct a historical decomposition by extrapolating the real price of oil from 2001.5 on the basis of the observed forecast surprises. All regressions are based on the full sample. The first panel of Figure 3 shows that $S_{1,t}$ cumulatively explains a substantial part of the increase in the real price of oil starting in mid-2004. It is not the only explanation, however. The upward pressure on the price is reinforced by the cumulative effect of $S_{2,t}$ starting in early 2004. Although these two news shocks are not uncorrelated, their correlation is fairly low. Thus, to a first approximation, we can add the two fitted values to obtain a crude estimate of the combined effect of global demand pressures on the real price of oil. The third panel shows that $S_{1,t}$ and $S_{2,t}$ jointly explain the bulk of the increase and decline of the real price of oil between 2002.1 and 2008.12, underscoring the explanatory power of demand shifts for the real price of oil. While these results have to be taken with a grain of salt because the two shock series are weakly correlated, they are suggestive.

4. Interpreting the Real GDP Forecast News

The impulse response analysis and historical decompositions in section 3 suggest that the news shocks identified in this paper trigger dynamics in the real price of oil of the type one would expect in response to demand shocks driven by the global business cycle. An interesting question is how the revisions in professional real GDP forecasts relate to the innovations to the demand

[14] Although it might be instructive to evaluate these estimates in light of a fully specified DSGE model, even the most recent generation of open economy DSGE models with endogenous oil prices does not include the type of global demand shock underlying the results in Figure 2 (see, e.g., Bodenstein, Erceg, and Guerrieri 2007).

for industrial commodities including crude oil identified in Kilian's (2009a) structural VAR model of the global crude oil market. There are three potential links. One is that the real GDP forecasters are wrong because they are surprised by contemporaneous demand innovations in global commodity markets, which in turn instantaneously move domestic real GDP. That view seems implausible, because a positive demand shock driven by orders for industrial commodities abroad, for example, would take time to stimulate the domestic economy. Indeed, the contemporaneous correlation of the two shock series is only weakly positive.

A more plausible view is that domestic real GDP responds with a delay to demand innovations in global industrial commodity markets. In particular, shipments of industrial commodities (which is in essence what the real activity index of Kilian (2009a) measures) can be interpreted as an indication of future increases in domestic output and investment. If that view were correct, one would expect real GDP forecast surprises to lag behind the aggregate demand shock series of Kilian (2009a).

A third possibility is that upward revisions of real GDP forecasts, especially in India and China, are a signal of a booming world economy and prompt additional demand for industrial raw materials, metals, and crude oil. If that view were correct, one would expect real GDP forecast revisions for emerging Asia to lead innovations to the global demand for industrial commodities.

Of course, these explanations are not mutually exclusive. Table 4 shows the cross-autocorrelation function of forecast surprises in India and China combined and of the VAR based global aggregate demand shock. It is striking that the highest correlations are fairly evenly distributed over adjacent leads and lags rather than being concentrated on one period only. No individual correlation is higher than 0.3. While the two shocks occur at approximately the same

time, which helps explain the broad similarity in the estimated responses, they are clearly not identical. This evidence is consistent with the explanations above being complementary rather than mutually exclusive. Given the contemporaneous correlation of only 0.14, the first explanation does not seem to play a major role. The highest correlations are clustered at $t+2$ (consistent with the second explanation) and at $t-1$ through $t-3$ (consistent with the third explanation). Based on the peak of the cross-autocorrelation function we conclude that typically professional forecast surprises lead innovations to the global demand for industrial commodities by one month, although the slow decline of the cross-autocorrelation function at longer lags suggests that the feedback to global industrial commodity markets is gradual and builds over time.

While this evidence is not dispositive about causal links between forecast news and global demand shocks, it is consistent with the interpretation that professional real GDP forecast revisions, while not being the same as global aggregate demand shocks in industrial commodity markets, are noisy signals of unexpectedly rising demand in industrial commodity markets including the market for crude oil. An interesting question is where these real GDP forecast revisions originate especially during 2003-2008. One interpretation is that they could reflect learning by forecasters about persistent changes in productivity resulting in a shift in potential real GDP (see, e.g., Egde, Laubach and Williams (2007) for a related model). That interpretation would miss that point that by all accounts demand for oil and other industrial commodities in countries such as China grew much faster than productivity growth could possibly have grown. Rather than an improvement in aggregate productivity holding fixed capital and labor inputs, that demand surge is likely to have reflected the accumulation of capital, notably the building of infrastructure and the structural transformation of the Chinese economy toward a more modern

energy-dependent society. That type of transformation is consistent with growth in the consumption of steel, cement, oil and other raw materials far in excess of aggregate real GDP growth rates. Put differently, if there was learning it was not so much about real GDP, but about the future demand for industrial raw materials, metals, and crude oil.

This interpretation also highlights that a given forecast revision in a country such as China would be expected to have a disproportionate effect on the demand for industrial commodities including oil compared with a forecast revision of the same magnitude for industrialized countries such as the United States, and underscores that aggregating forecast surprises across countries makes sense only if the countries in question have similar propensities to import industrial commodities. For that reason, it would not make sense to base the results in Figure 3 on a GDP-weighted aggregate of forecast surprises in the United States, Germany, Japan, India and China, for example.

5. Conclusion

Many explanations have been proposed for the surge in the real price of crude oil after 2003 including speculation in oil futures and spot markets, adverse oil supply shocks, deliberate restrictions on OPEC crude oil production, and shifts in global real economic activity. Understanding the causes of that increase is important for understanding the macroeconomic effects of oil price shocks and for the design of policy responses (see, e.g., Kilian 2008). This paper added to a growing body of evidence that the latest surge in the real price of oil is explained primarily by rising global demand for industrial commodities driven by unexpected economic growth. Our analysis differed from earlier work in that we utilized a direct measure of demand shocks (based on revisions of professional real GDP growth forecasts) rather than inferring demand shocks from an econometric model. We showed (1) that recent forecast

surprises were associated primarily with unexpected growth in emerging economies (in conjunction with much smaller positive GDP-weighted forecast surprises in the major industrialized economies), (2) that markets were repeatedly surprised by the strength of this growth, (3) that these surprises were associated with a hump-shaped response of the real price of oil which reaches its peak after 12-16 months, and (4) that news about global growth predict much of the surge in the real price of oil from mid-2003 until mid-2008 and much of its subsequent decline.

References

Alquist, R., and L. Kilian (2010), "What Do We Learn from the Price of Crude Oil Futures?" forthcoming: *Journal of Applied Econometrics*.

Andersen, T.G., and T. Bollerslev (1998), "DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies," *Journal of Finance*, 53, 219-265.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and C. Vega (2003), "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review*, 93, 38-62.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and C. Vega (2007), "Real-Time Price Discovery in Stock, Bond, and Foreign Exchange Markets," *Journal of International Economics*, 73, 251-277

Arsenau, D., M. Beechey, and R.J. Vigfusson, R.J. (2008), "The Response of Oil Prices to Inventory Surprises," mimeo, Federal Reserve Board.

Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.

Berkowitz, J., Birgean, I., and L. Kilian (1999), "On the Finite-Sample Accuracy of Nonparametric Resampling Algorithms for Economic Time Series," in T.B. Fomby and C. Hill (eds.): *Advances in Econometrics: Applying Kernel and Nonparametric Estimation to Economic Topics*, vol. 14, JAI Press, Connecticut, 77-107.

Bodenstein, M., Erceg, C.J., and L. Guerrieri (2007), "Oil Shocks and U.S. External Adjustment," mimeo, Federal Reserve Board.

Considine, T.J. (2006), “Is the Strategic Petroleum Reserve our Ace in the Hole?” *Energy Journal*, 27, 91-112.

Edge, R., Laubach, T., and J.C. Williams (2007), “Learning and Shifts in Long-Run Productivity Growth,” *Journal of Monetary Economics*, 54, 2421-2438.

Faust, J., J. Rogers, S.-Y. Wang, and J. Wright (2007), "The High Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements," *Journal of Monetary Economics*, 54, 1051-1068.

Frenkel, J.A., and A.K. Rose (2009), “Determinants of Agricultural and Mineral Commodity Prices,” mimeo, UC Berkeley.

Hamilton, J.D. (2009a), “Understanding Crude Oil Prices,” *Energy Journal*, 30, 179-206.

Hamilton, J.D. (2009b), “Causes and Consequences of the Oil Shock of 2007-08,” *Brookings Papers on Economic Activity*, 1, Spring, 215-261.

Kilian, L. (2008), “The Economic Effects of Energy Price Shocks,” *Journal of Economic Literature*, 46, 871-909.

Kilian, L. (2009a), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99, 1053-1069.

Kilian, L. (2009b), “Comment on ‘Causes and Consequences of the Oil Shock of 2007-08’ by James D. Hamilton,” *Brookings Papers on Economic Activity*, 1, Spring, 267-278.

Kilian, L. (2010), “Explaining Fluctuations in Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market,” forthcoming: *Energy Journal*.

Kilian, L., and C. Park (2009), “The Impact of Oil Price Shocks on the U.S. Stock Market,” *International Economic Review*, 50, 1267-1287.

Kilian, L., Rebucci, A., and N. Spatafora (2009), "Oil Shocks and External Balances," *Journal of International Economics*, 77, 181-194.

Kilian, L., and C. Vega (2010), "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices," forthcoming: *Review of Economics and Statistics*.

Smith, J. (2009), "World Oil: Market or Mayhem?" *Journal of Economic Perspectives*, 23, 145-164.

Table 1: Average Forecast Surprises (Percentage Points)

	2000.12-2003.5	2003.6-2008.6	2008.7-2008.12
USA	-0.05	0.02	-0.08
Germany	-0.12	0.00	-0.33
Japan	-0.10	0.08	-0.27
Brazil	-0.10	0.03	0.07
Russia	0.06	0.12	-0.42
India	-0.06	0.03	-0.17
China	-0.04	0.12	-0.17

SOURCE: Computations of the authors based on successive annual forecasts of real GDP growth reported by the *Economist Intelligence Unit*.

Table 2: Shares in World Real GDP (Percent)

	1996	2002	2009
USA	21.1	20.7	19.9
Germany	4.9	4.3	3.6
Japan	7.6	6.4	4.9
USA+Germany+Japan	33.6	32.2	28.3
OECD	58.4	55.4	51.0
Brazil	3.0	2.6	2.3
Russia	3.0	3.3	3.8
India	5.3	6.1	7.2
China	9.2	12.3	15.9
China+India	14.4	18.4	23.1

SOURCE: Computations of the authors based on the purchasing power adjusted real GDP data in the Penn World Table 6.2. The shares for 2009 have been obtained by linear extrapolation.

Table 3: GDP-Weighted Average Forecast Surprises (Percentage Points)

	2000.12-2003.5	2003.6-2008.6	2008.7-2008.12
USA	-0.010	0.005	-0.017
China+India	-0.006	0.017	-0.038
Germany+Japan	-0.014	0.006	-0.025
China+India+USA+Germany+Japan	-0.030	0.028	-0.080
OECD	-0.058	0.010	-0.085

SOURCE: Computations of the authors based on successive annual forecasts of real GDP growth reported by the *Economist Intelligence Unit*. The weights are the shares in PPP-weighted world real GDP.

Table 4: Cross-Autocorrelation Function for Real GDP Forecast Surprises in India and China Combined with respect to Global Aggregate Demand Shocks in Industrial Commodity Markets as Identified in the VAR Analysis of Kilian (2009a)

t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4
0.16	0.21	0.22	**0.28**	0.14	0.04	0.19	-0.10	0.09

SOURCE: The peak of the cross-autocorrelation function is shown in boldface, indicating that forecast surprises in India and China lead global aggregate demand shocks. The data sources are a suitably updated version of Kilian (2009a) and Figure 1.

Figure 1: Weighted Real GDP Growth Forecast Surprises: 2000.12-2008.12

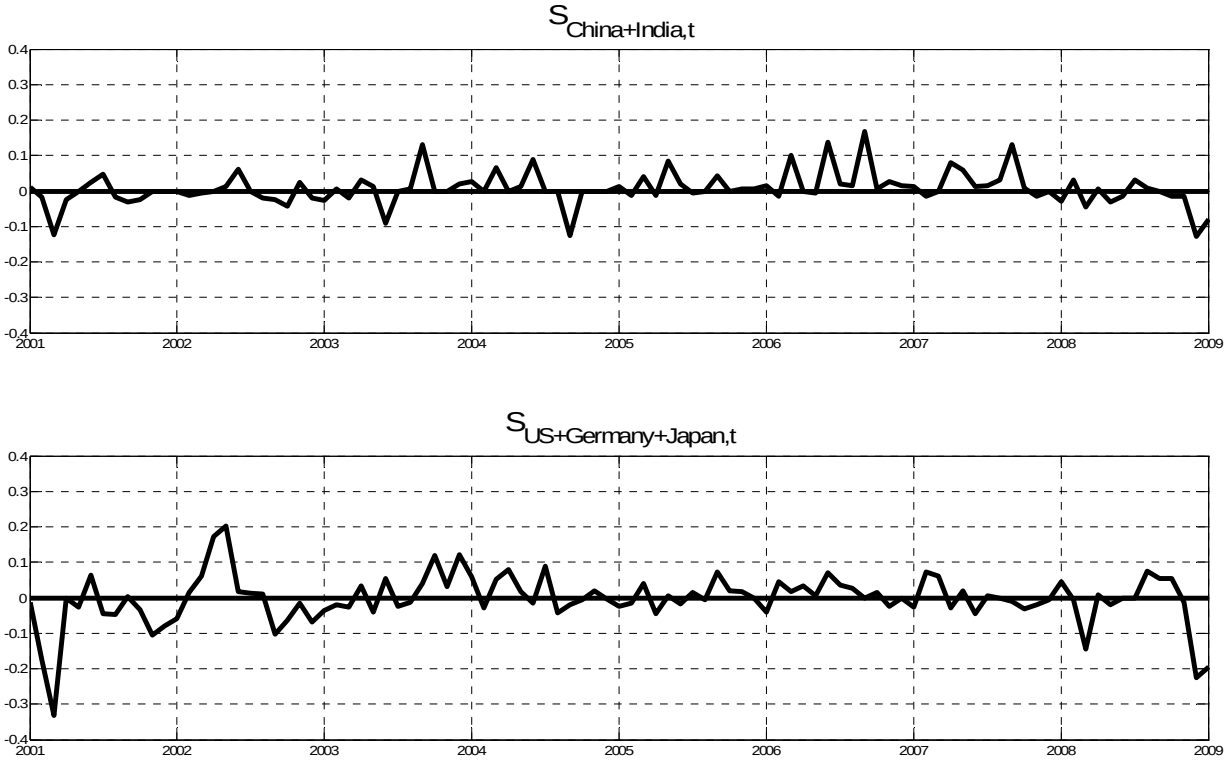

SOURCE: EIU revisions of annual real GDP growth forecasts for each country, constructed by the authors as described in the text and weighted by each country's share in purchasing power adjusted world real GDP, computed from data in the Penn World Table.

Figure 2a: Responses of the Real Price of Oil to Real GDP Forecast Surprises
Sample Restricted to Months with Nonzero Forecast Surprises
(with 1-Std. Error Bands)

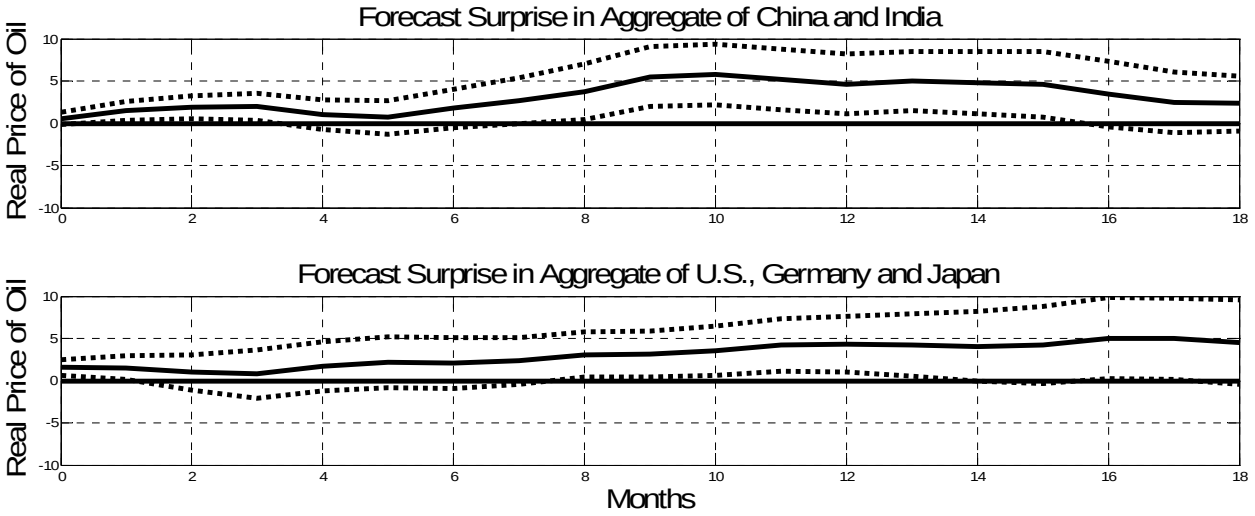

Source: Estimates based on regression model (1)**.**

Figure 2b: Responses of the Real Price of Oil to Real GDP Forecast Surprises: Full Sample
(with 1-Std. Error Bands)

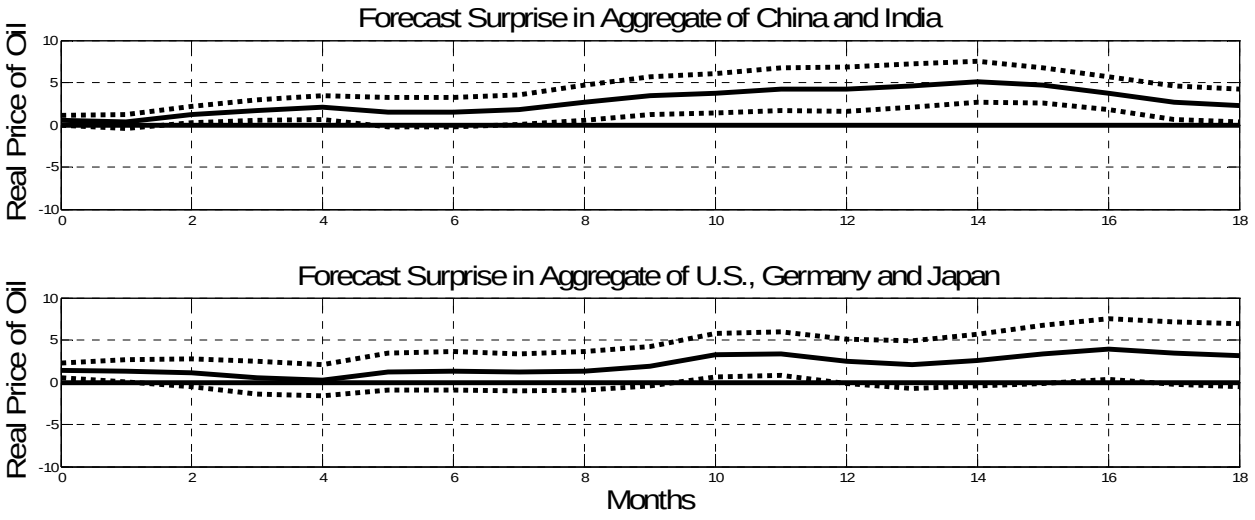

Source: Estimates based on regression model (1)**.**

Figure 2c: VAR Response to Global Aggregate Demand Shock (with 1-Std. Error Bands)

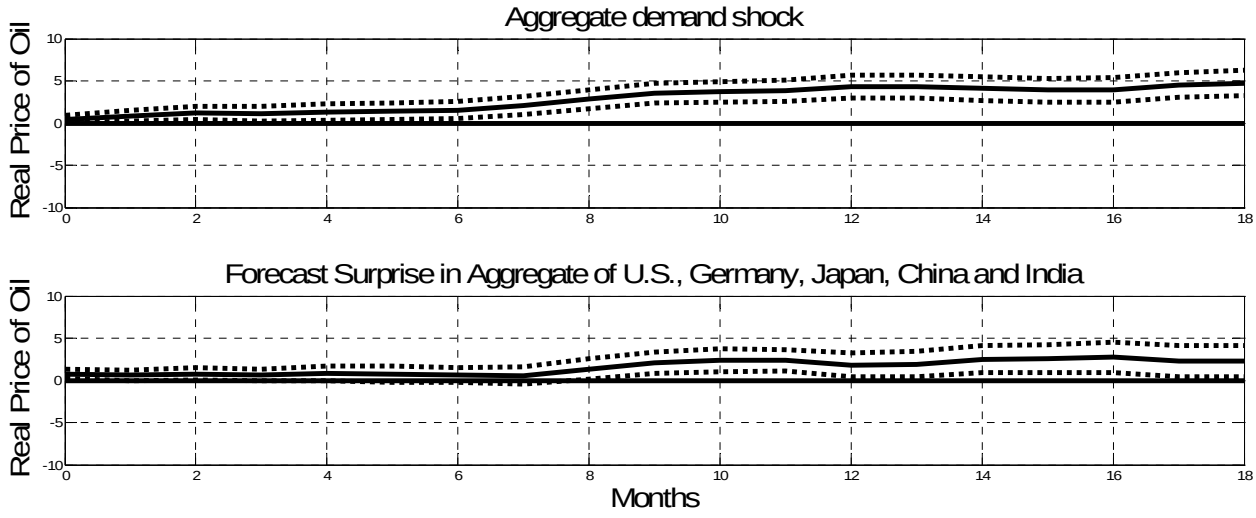

Source: Estimates based on regression model (1) and the VAR model in Kilian (2008, 2009).

Figure 3: Fitted Value of the Real Price of Oil

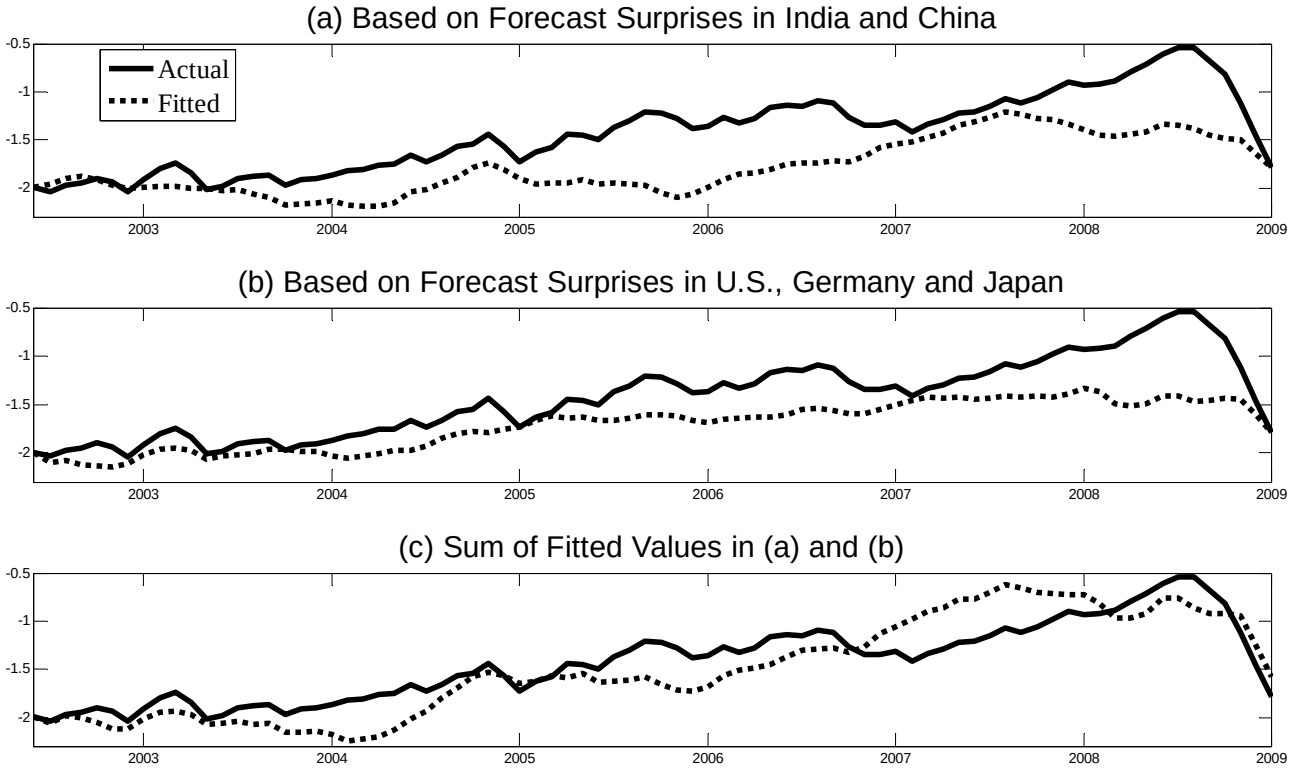

NOTES: Fitted values based on regression model (1) including periods in which forecast surprises are zero. Projection of the real price of oil conditional on its value in 2001.5.