

In The Americas:

1200, McGill College Av.
Suite 1000
Montreal, Quebec
Canada H3B 4G7

T + 1 514.878.2377

Contact details:

Christopher Cook
VP, Macro Solutions
cook@numeraanalytics.com

The Numera Approach

The purpose of this report is to illustrate our approach to macroeconomic analysis. We focus on crude oil prices for two reasons. First, future uncertainty is exceptionally high, emphasizing the need for probability forecasts. Second, oil price shocks are an important source of business cycle fluctuations, allowing for an especially rich assessment of macroeconomic dynamics.

Empirical macro modelling may be classified into two broad categories. On the one hand, there are models aimed at analyzing specific developments, such as the impact of policy changes on market outcomes. At their best, these models incorporate theoretical insights without compromising statistical fit. The first two sections of this report illustrate how these techniques may be used to support investment strategy. At the other end of the spectrum, we find models aimed at prediction. Unlike structural models, their quality depends on their track record. Despite multiple sources of uncertainty, practitioners rarely provide detailed information on forecast imprecision. By evaluating the full range of possibilities, sections 3 and 4 demonstrate how we use probability forecasts to monitor economic and financial risks while striving to minimize errors in prediction.

1: Drivers decompositions

Policy institutions typically use dynamic simultaneous equation models to study the impact of particular shocks or policies on the macroeconomy. In practice, of course, economic aggregates depend on a variety of interrelated factors, whose importance may vary over time. A natural, yet far less common application, is to use these models to disentangle the relative contribution of different factors to the actual evolution of a variable of interest. The approach is especially useful for strategic decisions, as it helps identify the forces driving demand, input costs, market and asset prices in specific industries.

We illustrate this procedure by analyzing historical fluctuations in the real (CPI-adjusted) WTI spot price. Like any global commodity, oil price changes are a consequence of shifting supply and demand conditions. Sudden movements receive ample media coverage, with industry pundits offering sometimes contentious explanations for their underlying cause¹. Although energy economists dissent on the relative importance of different transmission channels, most agree spot prices depend on fluctuations in physical supply, global demand for commodities and shifts in market sentiment (a broad term encompassing speculative changes in demand as well as uncertainty about the future).

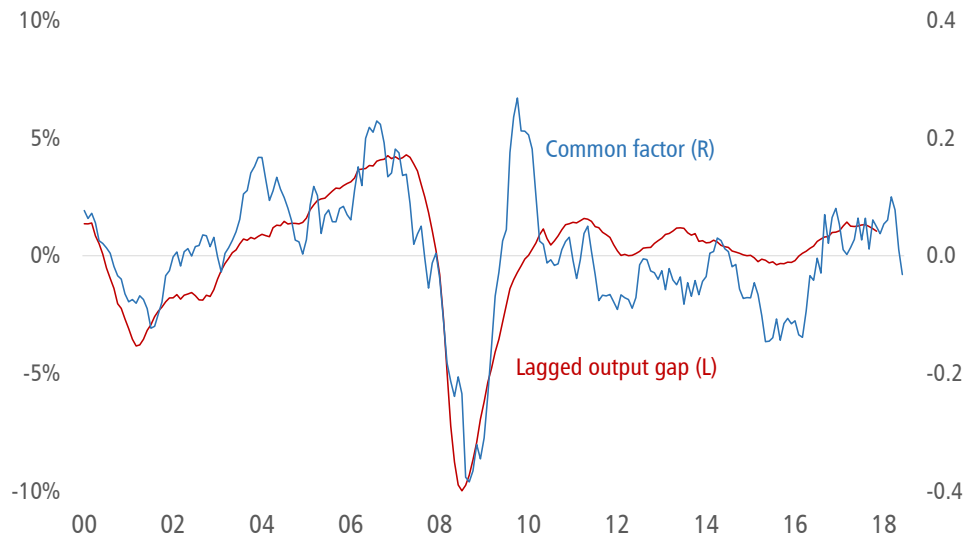
We model supply and demand conditions by employing a structural econometric model that allows for dynamic interactions between real oil prices, global economic activity, the trade-

¹ For example, see Story, L. (2008, May 21). "An oracle of oil predicts \$200-a-barrel crude". New York Times

weighted value of the US dollar, oil supply and above-ground crude oil inventories. The procedure is an extension to Kilian's (2009) seminal model for the global crude oil market. The two main differences is that our approach allows for time-varying cointegration, and estimates global real activity by extracting a latent factor from a panel of real commodity prices.

The technique is rooted on the premise that supply shocks are market-specific, so that common variations in commodity prices reflect cyclical fluctuations. The resulting diffusion index has the advantage of not relying on arbitrary weights (e.g. like industrial output), of not being distorted by idiosyncratic factors (e.g. like steel production), and of being available at a higher frequency, for a longer sample period and on a much more timely basis than alternative global indicators². The diffusion index is plotted in Figure 1 and compared against our measure of the global manufacturing output gap. Both indicators point towards much weaker global activity since 2011, with cyclical conditions only improving in 2017.

Figure 1: Alternative global demand measures (2000-2018)



Source: Numera Analytics

Once estimated, we can use the model to estimate the historical contribution of different supply and demand shocks to the real price of oil. Figure 2 plots year-on-year growth in real crude oil prices decomposed into five different components, so that the sum of each yields the observed price change at every point in time. Similarly, Figure 3 shows the cumulative breakdown during the first 8 months of 2018 (+31% in real terms). Because our model allows for second-round effects, each bar captures both direct and indirect channels of transmission (e.g. a weaker dollar may support demand for commodities, such that the resulting increase partly reflects an improvement in global activity).

Our results suggest changes in physical production have a much more limited effect on crude oil prices than demand-side factors. In fact, a variance decomposition over the period 1982-2018 reveals supply shocks only accounted for 5% of the variability in real crude oil prices. This does

²See Chiaie, Ferrara and Giannone 2017 for details on its estimation, and Kilian and Zhou 2018 for a comparison of global indicators.

not mean supply-side factors do not matter. Shifts in investor sentiment (primarily reflecting expectations about future supply disruptions) have been the main source of volatility over the past four decades, accounting for roughly 46% of the variability during this period.

Figure 2: Drivers decomposition: Real WTI (1991-2018 YOY%)

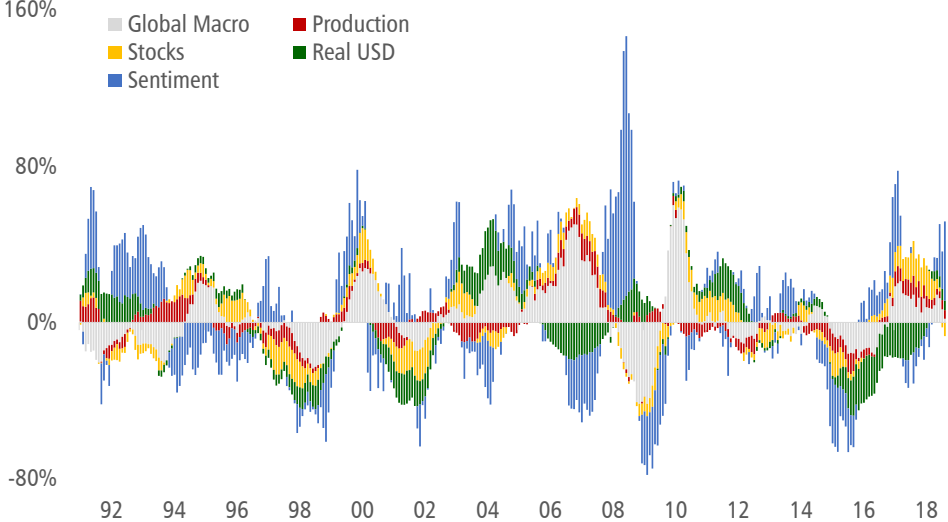
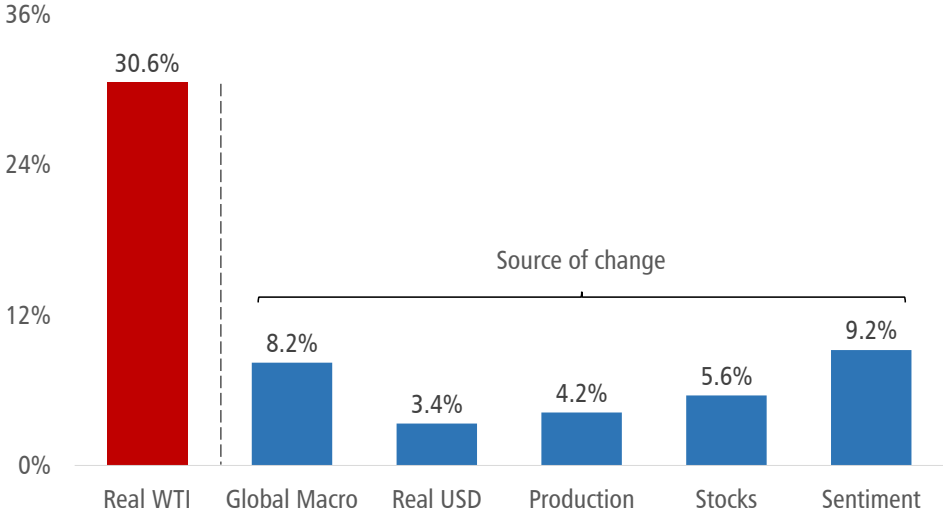


Figure 3: Drivers decomposition: Real WTI (8M 2018 YTD%)




Furthermore, production shocks may play an important role at specific points in time. For example, during the first eight months of 2018, restrained production by OPEC members and sharp declines in crude oil stocks in advanced economies contributed roughly one-third of the increase in real oil prices. A strong global economy and upbeat trader sentiment were the two most important sources of change over this period.

2: Impact assessment

Suppose you wanted to understand the effects of specific shocks or policies. For example, the impact of weaker growth in China on European manufacturing, the long-term consequences of Brexit on asset returns, or whether the speed, magnitude or duration of responses to Fed rate changes depend on the state of the economic cycle. We illustrate our approach to these types of questions by exploring the short-run implications of rising oil prices on the US economy. More specifically, we evaluate the response of selected US indicators to price increases triggered by: a) a stronger world economy, and b) speculative changes in demand.

Given the size of the US economy, cyclical fluctuations have important international ramifications. Rather than assuming the oil market to be unaffected by US developments, we augment our baseline specification with a core US model and allow for feedback between both blocks. The US block contains a series of long run theoretical restrictions (e.g. arbitrage conditions linking prices and returns over time) that allow for a clear interpretation of the results. We do not restrict the short-term dynamics, as theory is typically silent on the precise sequencing of economic outcomes.

Figure 4: 12-month cumulative responses to oil price increases



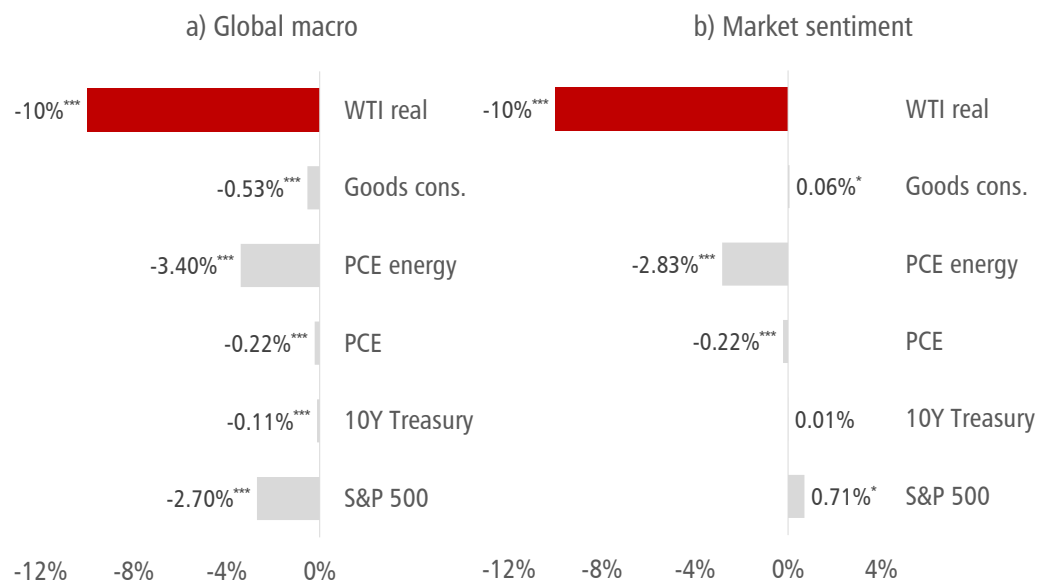
Note: Bars denote cumulative generalized impulse response functions (Pesaran and Shin 1998). Asterisks refer to statistical significance at the 90/95/99% confidence level. The p-values are based on 2000 moving block bootstrap (MBB) replications to control for volatility clustering in the residuals (see Brüggemann et al. 2016 for details on its implementation).

Figure 4 shows the cumulative response of selected economic indicators to a 10% increase in crude oil prices over a 12-month period. One may interpret these effects as how much each variable would deviate from its baseline projection if oil prices were subject to one-time shocks. The difference in results highlights the importance of analyzing different channels of transmission. If

the price increase arises from stronger global conditions, goods consumption rises 0.8% despite a 0.25% increase in consumer prices. Equity prices rise 4.3% above baseline, while the 10-year Treasury bond grows 21 basis points reflecting stronger growth prospects, higher inflation expectations and a portfolio rebalancing effect. The results are markedly different if triggered by speculation. In this case, higher consumer prices (0.17% above trend) cause an unambiguous loss of purchasing power so goods consumption falls 0.2%. Reduced growth prospects offset the higher inflationary expectations, leaving long-term yields roughly unchanged. Overall equity prices fall 1.2% due to the reduced earnings potential.

There are two aspects to the analysis worth emphasizing. Firstly, structural instability in some of the underlying relations means the effects may vary over time. For instance, the impact of a macro-driven shock on equity prices is stronger today than it was in the 1990's (about 3.4% for an equivalently sized shock). Secondly, there is no guarantee that oil price declines have the same but opposite effect on real activity. Figure 5 shows the impact of a 10% cumulative decline in oil prices arising from the same two sources. Notice that the multipliers are considerably smaller. Consider, for instance, the effects of a 10% drop triggered by an unexpected weakening of the global economy. The pass-through to consumer prices is about 20% lower than if the world economy were strengthening, evidence of downward price rigidities.

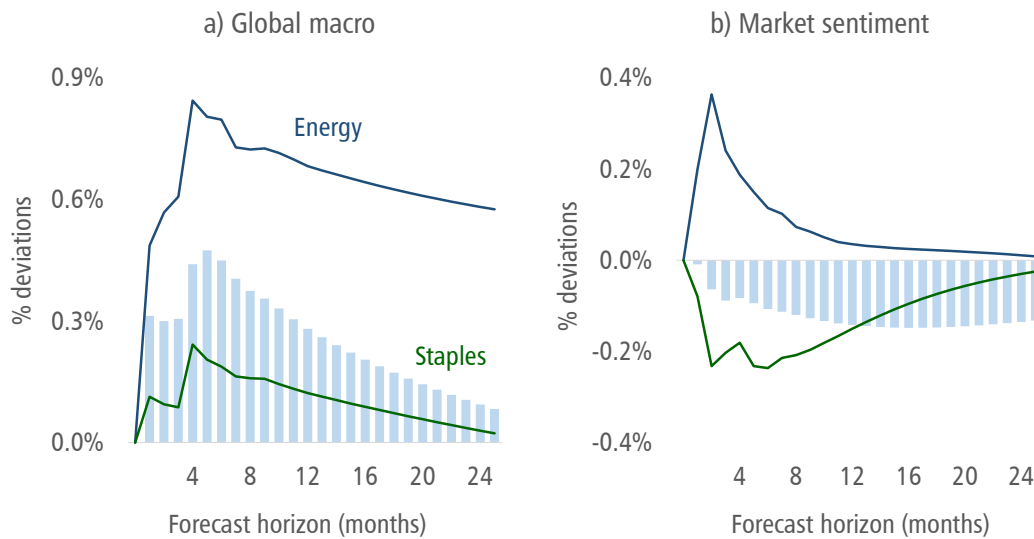
Figure 5: 12-month cumulative responses to oil price declines



Of course, we can extend the analysis to study sectoral implications. For instance, one may wish to explore the effect of higher oil prices across equity categories. Clearly, the exposure of energy companies (or suppliers to the mining industry) will be different from segments dependent on household spending. Figure 6 plots the time path to both oil market shocks over a 2 year span. Each line measures monthly deviations from baseline projections for the S&P 500 (light blue bars) and its energy and consumer staples components. Notice global macro shocks have a larger and more persistent effect, even causing a level shift in energy stocks

(the 12-month cumulative increase is 8.4%, with prices remaining well above baseline after 2 years). The impact on defensives is limited, with equity prices rising 1.8% after one year. The market sentiment shock has a transitory effect on both sub-components, with the direction of the response reversed. While staples contract 2.3% after 12 months, the average stock price of US energy companies grows 1.6% as higher oil and gasoline prices more than compensate for weaker expected demand for energy goods.

Figure 6: Sectoral responses to oil price increases



3: Forecast evaluation

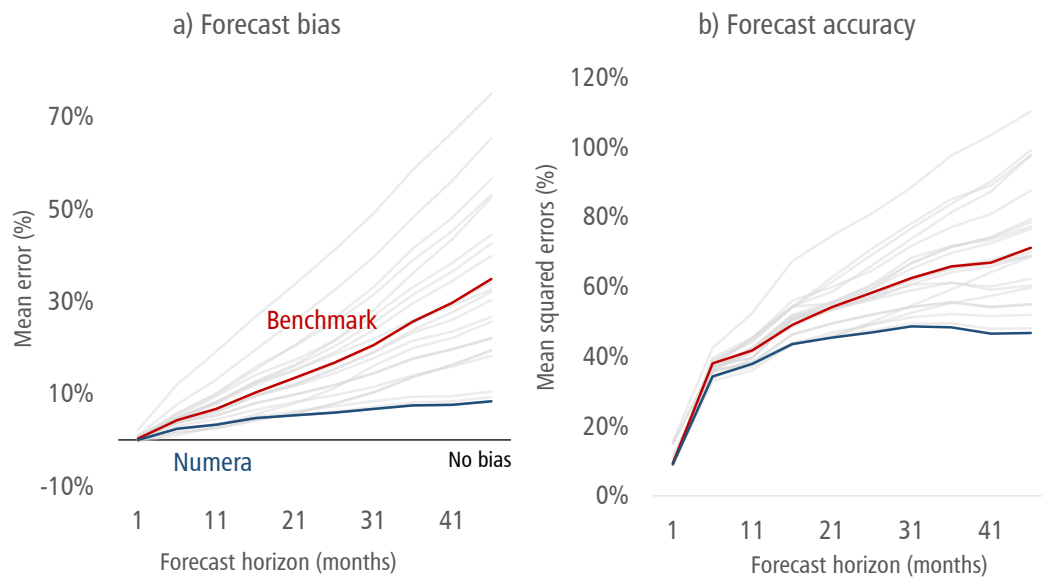
Forecast uncertainty stems from multiple sources, including model choice, inaccuracy in the value of coefficients, structural change, omitted predictor bias, and the inherent unpredictability of future shocks. In addition, while structural models may produce reliable projections, good in-sample fit does not generalize into the future. For example, it may be difficult to anticipate future developments in related series. Our approach to economic forecasting starts by testing and comparing the forecasting (out-of-sample) performance of a large selection of alternative models. This ensures the projections are based on models with limited bias, and a high degree of accuracy relative to forecasting benchmarks.

Potential models differ in their complexity and reliance on economic theory. While model choice varies widely across concepts, our most successful specifications usually incorporate novel insights from the forecasting literature, such as data shrinkage methods or dynamic model averaging. The exhaustiveness of the forecast evaluation stage cannot be overstated. Our **Global Macro Monitor**, for example, is the product of a search across 3000+ possible models for the different concepts featured in the report.

Figure 7 shows bias and accuracy statistics for real WTI prices over the 2005-2018 period. The left hand panel shows the average forecast errors from 1 month to 4 years ahead for a selection

of competing models. The mean error (panel A) measures by how much forecasts deviate from realized values, so that a positive figure signals upward bias. Panel B plots the average squared forecast errors (the forecast error variance), a commonly used measure of point forecast accuracy. In line with the commodity price forecasting literature, we set the benchmark to the no change forecast (see for example Alquist et al. 2011). Note that for the evaluation to be valid, the projections cannot be based on future information (e.g. taking the future observed value of the US dollar as given), as this would artificially lower the errors in prediction.

Figure 7: Forecast evaluation summary statistics (2005-2018)

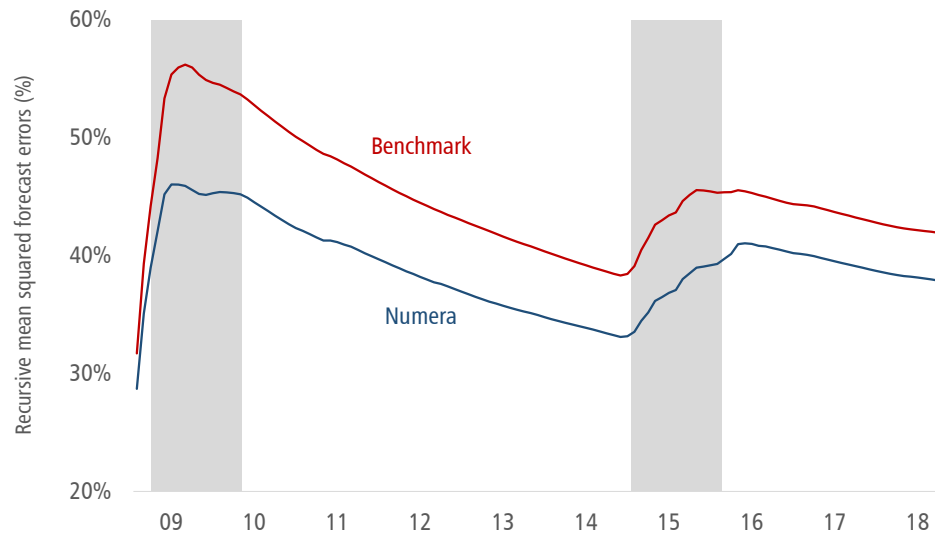


The blue lines correspond to our preferred specification ('Numera'). We employ a forecast combination approach that exploits the predictive content of individual models at different points in time. We update the weights dynamically based on their past forecasting performance. More specifically, we switch between 3 specifications, one of which is a modified version of the model employed in section 1. Compared to alternative methods, the technique exhibits limited upward bias and is generally more accurate than the benchmark. Although the predictive performance is similar in the short-run (no-change forecasts are very difficult to beat at short horizons), accuracy improves markedly after one year. At 36 months, the model beats the benchmark by 25%. The results are in line with Baumeister and Kilian's (2013) findings for the real US refiners' acquisition cost of oil.

One potential limitation with summary statistics is that they may hide differences in predictability across the evaluation period. In particular, low mean squared errors may simply reflect a handful of periods of good forecasting performance. One way to determine if the results are robust to the choice of forecast sample is to compute the squared prediction errors recursively. Figure 8 shows the results for one-year ahead projections. Unsurprisingly, predictability worsens during the Great Recession, with forecast errors almost doubling in size after H2/2009. Forecast uncertainty also increases in mid-2014, although the impact on the forecast accuracy measure

is less pronounced. The important thing, however, is that accuracy remains consistently higher for our chosen specification.

Figure 8: Recursive 1-year ahead squared forecast errors (%)



4: Risk monitoring

Even with well-specified forecast models, future shocks will necessarily cause projections to deviate from their outcome. Point forecasts offer an especially poor assessment of growth prospects for variables that are subject to multiple sources of uncertainty. For example, the 3-year ahead prediction errors for real oil prices are 15-20 times higher than the corresponding forecast errors on headline CPI measures for advanced economies.

In our view, probability forecasts offer a much more comprehensive assessment of future growth prospects. Unlike point forecasts, these are statements of the likelihood of specific events taking place³. As a result, they shift the emphasis from central outcomes to the monitoring and evaluation of macroeconomic and financial stress.

The difficulty in deriving density forecasts for economic indicators stems from the fact risk varies over the course of the business cycle. Thus, the size and shape of predictive distributions will be time varying. Figure 9 plots the conditional volatility on the 1-month ahead forecast errors on crude oil prices. Future uncertainty spikes during periods of economic and financial turmoil, but quickly reverts to trend. This has important implications for the measurement of downside risks.

Figure 10 plots the 3-month and 1-year ahead predictive densities for real crude oil prices as of September 2018. At one year out, not only does the distribution flatten (indicating greater uncertainty), but downside risks for oil firms increase substantially. This can be seen from the shape of the cumulative distributions, with the slope becoming flatter and tilted towards the left.

³ Rossi (2014) offers a good introduction to density forecasting in macroeconomics.

The densities allow us to evaluate specific scenarios. For instance, imagine shale oil producers can support oil drilling as long as the real price of crude exceeds \$50 (this is just an assumption). The probability of real oil prices falling below \$50 in 3 months time is about 2.3%, but the likelihood rises to 23.9% by Q3/2019. The same exercise, carried out in periods of high uncertainty, would deliver vastly different results, with the probability of extreme events increasing markedly even in the very short term.

Figure 9: Short-term volatility in crude oil prices (2006-2018)





Figure 10: 3 and 12-month ahead forecast distributions (%)



This special report was written by Joaquín Kritz Lara, head of Numera’s Macro Research practice. For details on our macro research offering, please contact Christopher Cook at ccook@numeraanalytics.com or +1.514.861.8848.

References

- Alquist, R., Kilian, L., and Vigfusson, R.J. (2011) "Forecasting the price of oil" CEPI Discussion Papers No 8388
- Chiaie, S., Ferrara L., and Giannone, D., (2017) "Common factors of commodity prices" ECB Working Paper Series No 2112
- Kilian, L. (2009). "Not all price shocks are alike: Disentangling demand and supply shocks in the crude oil market" American Economic Review 99-3, 1053-1069
- Kilian, L., and Zhou, X., (2018) "Modeling fluctuations in the global demand for commodities" Bank of Canada Staff Working Paper 2018-4
- Pesaran, M.H., and Shin, Y. (1998) "Generalized impulse response analysis in linear multivariate models", Economic Letters 58, 17-29
- Brüggemann, R., Jentsch, C., and Trenkler, C. (2016) "Inference in VARs with conditional heteroskedasticity of unknown form" Journal of Econometrics 191: 69-85
- Rossi, B. (2014). "Density forecasts in economics and policymaking" Els Opuscles del CREI, N37
- Baumeister, C., and Kilian, L., (2013). "Forecasting the real price of oil in a changing world: A forecast combination approach", Journal of Business and Economic Statistics 33(3), 338-351