



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2012-1

Real-Time Analysis of Oil Price Risks Using Forecast Scenarios

by Christiane Baumeister and Lutz Kilian

Bank of Canada Working Paper 2012-1

January 2012

Real-Time Analysis of Oil Price Risks Using Forecast Scenarios

by

Christiane Baumeister¹ and Lutz Kilian²

¹International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
cbaumeister@bankofcanada.ca

²Department of Economics
University of Michigan
Ann Arbor, MI 48109-1220
lkilian@umich.edu

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors.

No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

William Wu provided excellent research assistance. We thank Norm Swanson and numerous seminar participants for helpful discussions.

Abstract

Recently, there has been increased interest in real-time forecasts of the real price of crude oil. Standard oil price forecasts based on reduced-form regressions or based on oil futures prices do not allow consumers of forecasts to explore how much the forecast would change relative to the baseline forecast under alternative scenarios about future oil demand and oil supply conditions. Such scenario analysis is of central importance for end-users of oil price forecasts interested in evaluating the risks underlying these forecasts. We show how policy-relevant forecast scenarios can be constructed from recently proposed structural vector autoregressive models of the global oil market and how changes in the probability weights attached to these scenarios affect the upside and downside risks embodied in the baseline real-time oil price forecast. Such risk analysis helps forecast users understand what assumptions are driving the forecast. An application to real-time data for December 2010 illustrates the use of these tools in conjunction with reduced-form vector autoregressive forecasts of the real price of oil, the superior real-time forecast accuracy of which has recently been established.

JEL classification: Q43, C53, E32

Bank classification: Econometric and statistical methods; International topics

Résumé

On note un intérêt accru, ces derniers temps, pour la prévision en temps réel du prix réel du pétrole brut. Les prévisions usuelles de ce prix, obtenues avec des équations de forme réduite ou fondées sur les cours à terme de l'or noir, ne permettent pas d'explorer l'incidence de scénarios différents au sujet de l'évolution future de la demande et de l'offre de pétrole sur la prévision de référence. L'analyse de divers scénarios est d'une importance capitale pour les utilisateurs finaux des prévisions désireux d'évaluer les risques entachant ces dernières. Les auteurs montrent : 1) qu'on peut élaborer des scénarios pertinents au regard des politiques à l'aide de modèles vectoriels autorégressifs structurels du marché pétrolier mondial récemment proposés; 2) que les variations des probabilités associées à ces scénarios influent sur les risques à la hausse et à la baisse entourant la prévision de référence en temps réel du prix réel du pétrole. Cette analyse des risques aide les utilisateurs des prévisions à cerner les hypothèses déterminantes. Une application de ces outils aux données en temps réel de décembre 2010 illustre comment on peut conjuguer ceux-ci à des prévisions du prix réel de l'or noir issues de vecteurs autorégressifs de forme réduite, dont la qualité supérieure a récemment été établie.

Classification JEL : Q43, C53, E32

Classification de la Banque : Méthodes économétriques et statistiques; Questions internationales

1. Introduction

The real price of oil is one of the key variables in the model-based macroeconomic projections generated by central banks, private sector forecasters, and international organizations. Recent research has shown that the real price of oil is forecastable in real time at horizons up to one year (see Baumeister and Kilian 2011). In particular, suitably chosen reduced-form vector autoregressive (VAR) models that include the key variables relevant to the determination of the real price of oil (such as global oil production, global real economic activity, and above-ground crude oil inventories) enjoy superior out-of-sample forecast accuracy. These VAR models have lower real-time mean-squared prediction error (MSPE) as well as higher directional accuracy than the no-change forecast, in particular. They also are more accurate than forecasts based on oil futures prices and forecasts based on univariate reduced-form models of the real price of oil.

Real-time forecast accuracy, however, is but one criterion relevant to the end-users of oil price forecasts. From the point of view of policymakers as well as private sector forecasters, forecasts based on reduced-form regressions (much like the corresponding forecasts based on oil futures prices) are of limited usefulness because they do not provide any economic insight into what is driving the forecast. In particular, reduced-form forecasts do not allow us to explore how much the oil price forecast would change relative to the unconditional baseline forecast under alternative scenarios about future oil demand and oil supply conditions. Such scenario analysis is of central importance for anyone interested in assessing the risks underlying oil price forecasts. The construction of *forecast scenarios* (also sometimes referred to as *conditional forecasts*) requires the use of structural models (see Waggoner and Zha 1999). In this paper, we illustrate how policy-relevant forecast scenarios about future oil demand and oil supply conditions can be constructed from recently proposed structural VAR models of the global oil market. Our analysis builds on the structural model proposed in Kilian and Murphy (2010), which also motivated the choice of the variables for the reduced-form VAR models used in the real-time forecasting analysis of Baumeister and Kilian (2011).

An application to real-time data for December 2010 illustrates the use of these tools in conjunction with reduced-form vector autoregressive forecasts of the real price of oil. We demonstrate that our reduced-form VAR forecast correctly predicted the pattern of the evolution of the real price of oil in 2011 to date, whereas conventional forecasts based on the random walk model or based on oil futures prices did not. We then illustrate how alternative forecast scenarios could have generated very different outcomes. The sensitivity of the real-time forecast of the real price of oil to alternative assumptions about future oil demand and oil supply conditions can be captured by probability-weighted predictive densities and summarized by formal risk measures. Risk analysis along these lines allows end-users of oil price forecasts to explore by how much and at what forecast horizon downside and upside risks change as a function of the probability weights attached to different scenarios, and facilitates an understanding of the tradeoffs between alternative assumptions.

The remainder of the paper is organized as follows. In section 2 we review the reduced-form model and structural model of the oil market underlying our analysis. Section 3 illustrates how the structural moving average representation of this model may be used to understand historical fluctuations in the real price of oil. Historical decompositions highlighting the cumulative contribution of different types of oil demand and oil supply shocks to the evolution of the real price of oil are useful in understanding the past. They are also invaluable in gauging what future oil supply and oil demand shock sequences may look like and help forecast users develop realistic forecast scenarios. In section 4, we discuss how the structural moving average representation of the model can be iterated forward in time to construct forecast scenarios. Section 5 provides several examples of such scenarios based on historical precedent as well as hypothetical assumptions. In section 6, we show how real-time forecasts of the real price of oil may be augmented by forecast scenarios to convey the sensitivity of the baseline forecast to changes in assumptions about future oil demand and oil supply shocks. Section 7 demonstrates how these forecasting scenarios may be weighted by their respective probabilities, allowing the construction of a probability-weighted predictive density for the real price of oil, which can be evaluated using standard risk measures. The concluding remarks are in section 8.

2. A Structural Model of the Determination of the Real Price of Crude Oil

Standard forecasting models are selected to produce low MSPE forecasts or to have high directional accuracy, with little regard to the underlying economic structure. One important limitation of such forecasting methods from the point of view of many end-users is that they do not shed light on why the real price of oil has moved in the recent past or is expected to move in the future. An equally important limitation is that such forecasting methods do not convey how sensitive the real oil price forecast is to hypothetical events in the global market for crude oil. Forecasts which condition on such hypothetical events are referred to as *forecast scenarios*. For example, a user may be interested in how an oil production shortfall would affect the forecast of the real price of oil. Similarly, one may be interested in exploring the possible consequences of civil unrest in the Middle East, or in exploring how much a period of unexpectedly low global demand for crude oil caused by a global recession would lower the real price of oil. The construction of such forecast scenarios requires the use of structural econometric models.

Structural models of the global market for crude oil have recently been developed by Kilian (2009), Kilian and Murphy (2010, 2011), and Baumeister and Peersman (2011), among others. In this paper, we build on the structural vector autoregressive model proposed in Kilian and Murphy (2010). This model was designed to help us distinguish, in particular, between unexpected oil production shortfalls, unexpected changes in the global demand for crude oil associated with fluctuations in the global business cycle, and shocks to the demand for above-ground crude oil inventories driven by speculative motives not captured by flow demand and flow supply shocks. If speculation were conducted by oil producers rather than oil consumers, with oil producers deliberately delaying the production of crude oil in anticipation of higher prices (causing an accumulation of below-ground oil inventories), the model by construction would treat such shocks as negative oil supply shocks. The structural model also allows for an additional residual demand shock designed to capture additional idiosyncratic demand shocks. The latter oil demand shocks have no systematic effect on the evolution of the real price of oil and can be safely ignored for the purpose of this paper. Changes in the composition of these structural shocks help explain why

conventional regressions of macroeconomic aggregates on the price of oil tend to be unstable. They also are potentially important in interpreting forecasts of the real price of oil.

2.1. Specification of the Reduced-Form Model

The reduced-form representation of the Kilian and Murphy (2010) model corresponds to the four-variable VAR model which has been shown in Baumeister and Kilian (2011) to have superior real-time forecast accuracy compared with the no-change forecast benchmark as well as other forecasting models. It includes the percent change in global crude oil production, an index of fluctuations in the global business cycle developed in Kilian (2009), the real price of crude oil, and the change in above-ground crude oil inventories.

The nominal price of oil in the global market is approximated by the U.S. refiners' acquisition cost for crude oil imports (see Alquist, Kilian, and Vigfusson 2011). The real price of oil is constructed by deflating the nominal dollar price of oil by the seasonally adjusted U.S. consumer price index for all urban consumers. World oil production is expressed in growth rates. Following Kilian and Murphy (2010), a proxy for the change in world crude oil inventories is constructed by scaling U.S. crude oil inventories by the ratio of OECD over U.S. petroleum inventories. This approximation is required because there are no monthly crude oil inventory data for other countries. Finally, as is standard in the literature, the measure of global real activity is constructed by cumulating the growth rate of the index of nominal shipping rates; the resulting nominal index is deflated by the U.S. consumer price index, and a linear deterministic trend representing increasing returns to scale in ocean shipping is removed from the real index (see, e.g., Kilian 2009). This index is designed to capture business cycle fluctuations in global industrial commodity markets. In constructing the real-time version of the index of global real activity, the linear deterministic trend is recursively re-estimated in real time. For a detailed discussion of the data sources the reader is referred to Baumeister and Kilian (2011).

The model explicitly embeds both stock and flow interpretations of the determination of the real spot price of oil. The inclusion of crude oil inventories allows the model to capture forward-looking

behavior based on data not observable to the econometrician. It also allows us to exclude data on the oil futures spread. The model exploits the fact that the spot and futures markets for crude oil are linked by an arbitrage condition (see Alquist and Kilian 2010). Thus, any speculation taking place in the oil futures market implies a shift in inventory demand in the spot market by construction. This fact allows us to abstract from the oil futures market altogether. Indeed, this theoretical result is consistent with the fact that the oil futures spread does not Granger-cause the four variables included in the model. This is an important test of the validity of the identifying assumptions, as discussed in Giannone and Reichlin (2006). It is important to stress that the model based on oil inventories remains well specified even in the absence of an oil futures market (which only emerged in the mid-1980s). This fact enables us to estimate the model based on data back to January 1973, which allows one to verify that the model estimates are also consistent with extraneous evidence on speculative activities in oil markets, for example, in 1979 (see, e.g., Yergin (1992), p. 687).

When conducting structural analysis based on the full sample, the model is estimated by least squares using 24 lags, as in Kilian and Murphy (2010). In recursive out-of-sample forecasting based on shorter samples, Baumeister and Kilian (2011) show that a more parsimonious specification involving 12 lags or alternatively the use of Bayesian shrinkage methods is preferred. In our real-time forecasting application in section 6, the choice of the lag order is less important given the much longer estimation sample available by 2010.¹² The forecasts implied by the VAR(24) and VAR(12) models are so similar that we use the VAR(24) model throughout this paper.

2.2. Identification of the Structural Shocks

The structural oil demand and oil supply shocks in this model are jointly identified based on a combination of sign restrictions on the structural impulse responses and bounds on impact price elasticities of oil demand and oil supply. The sign restrictions can be motivated based on standard economic theory. Specifically, the model imposes that unexpected disruptions of the flow supply of oil, conditional on past data, shift the oil supply curve along the oil demand curve, causing the price to

increase and the production of oil to fall on impact. The increase in the real price of oil is also associated with a fall in global real activity on impact. The sign of the impact effect on inventory accumulation cannot be pinned down a priori and is left unrestricted. On the other hand, in response to a positive flow demand shock, represented as a shift of the oil demand curve along the supply curve conditional on past data, price and production must increase on impact, as must global real activity. The sign of the impact response of inventories again is unrestricted. Finally, an unexpected shift in speculative demand all else equal implies in equilibrium positive inventory accumulation as well as an increase in the real price of oil on impact (see Alquist and Kilian 2010). The inventory accumulation requires an increase in oil production and a decline in oil consumption, associated with a decline in global real activity.

These sign restrictions on the impact responses of the model variables are further strengthened by bounding the impact price elasticities of oil supply and oil demand. These elasticities can be expressed as ratios of the impact responses of the quantity and price of oil. We impose an upper bound on the price elasticity of oil supply of 0.025, consistent with the consensus view in the literature that this elasticity is near zero in the short run (see, e.g., Hamilton 2009; Kilian 2009; Kellogg 2011). This bound could be relaxed somewhat without changing the substantive results. We also impose that the impact price elasticity of oil demand is bounded above by zero (implying that oil demand falls when the real price of oil rises) and is bounded below by the long-run price elasticity of oil demand, which can be inferred from independent cross-sectional estimates.

Finally, we impose that an unexpected flow supply disruption is associated with a positive response of the real price of oil and a negative response of oil production and global real activity for the first year. This latter assumption helps rule out models that are inconsistent with conventional views of the effects of flow supply shocks.

For the practical implementation of this approach to identification the reader is referred to Kilian and Murphy (2010). Although the Kilian and Murphy model is only set-identified, all structural models that satisfy the identifying restrictions can be shown to be quite similar, allowing us to focus on one such model with little loss of generality. We follow Kilian and Murphy in selecting the model with a price

elasticity of oil demand closest to the posterior median of that elasticity.

3. Historical Decompositions

Before constructing forecast scenarios, it is essential to understand what has been driving the real price of oil in the past. The point of studying historical decompositions of the real price of oil is not only to understand the determinants of historical oil price fluctuations, but also to guide us in developing realistic forecast scenarios on the basis of structural shock sequences observed during selected historical episodes of interest.

For the purpose of constructing the historical decompositions, we estimate the structural oil market model of Kilian and Murphy (2010) on ex-post revised data for 1973.1 to 2010.6. We exclude the more recent data which, as of the time of the real-time forecast discussed in section 6, were still incomplete and subject to revisions (see Baumeister and Kilian 2011). Under the maintained assumption of stationarity, the structural moving average representation of the estimated model allows us to decompose historical fluctuations in the real price of oil into orthogonal components corresponding to different oil demand and oil supply shocks. Let

$$y_t = \sum_{i=0}^{\infty} \Theta_i w_{t-i} \approx \sum_{i=0}^{n-1} \Theta_i w_{t-i}, \quad (1)$$

where y_t refers to the current observations of the data, Θ_i denotes the matrix of structural impulse responses at lag $i = 0, 1, 2, \dots$, and w_t denotes the vector of mutually uncorrelated structural shocks (see Lütkepohl 2005, chapter 3). The deterministic regressors have been omitted for expository purposes. In practice, Θ_i and w_t may be estimated consistently from the data. The approximation (1) improves over time. It is customary to discard observations at the beginning of the sample, for which the approximation tends to be poor.

The historical decomposition in Figure 1 shows the relative contribution of each oil demand and oil supply shock during key historical episodes. Each panel in Figure 1 plots the cumulative effect on the

real price of oil of one structural shock, while turning off all other shocks. The vertical bars correspond to key events in the global oil market including the Iranian Revolution in late 1978, the outbreak of the Iran-Iraq War in late 1980, the collapse of OPEC at the end of 1985, the August 1990 invasion of Kuwait, the Asian crisis of 1997, the Venezuelan oil supply shock of late 2002 (immediately followed by the Iraq War of early 2003), and the global financial crisis of mid-2008. These events provide useful information for the construction of the forecast scenarios, as discussed in section 5.

Each historical episode is characterized by a different mix of structural shocks. For example, the 1979 surge in the real price of oil was jointly caused by strong unexpected flow demand, associated with a boom in the world economy and in industrial commodity markets, and by a surge in speculative demand starting in May 1979 (with little contribution from the unanticipated oil supply disruptions associated with the Iranian Revolution earlier that year). In contrast, the decline in the real price of oil in 1998, for example, can be attributed primarily to an unexpected decline in flow demand for crude oil following the Asian crisis.

These and other episodes have been discussed at length in Kilian and Murphy (2010) except for the aftermath of the financial crisis. Figure 1 illustrates that much of the surge in the real price of oil between 2003 and mid-2008 was caused by repeated unexpected increases in the global business cycle as opposed to positive speculative demand shocks or negative flow supply shocks. As financial markets collapsed in the second half of 2008, so did global real activity and hence demand for industrial commodities such as crude oil. Essentially orders and shipping of industrial commodities ceased. The collapse in global demand far exceeded the corresponding decline in real GDP and industrial production because global shipping of industrial commodities was effectively suspended, as policymakers scrambled to restore financial markets. In early 2009, global demand recovered as quickly as it had collapsed, when confidence was restored. The cumulative effect of flow demand shocks reached a peak in January of 2010 that is only somewhat lower than the all-time high of mid-2008. In the first half of 2010 these demand pressures eased, however, and the cumulative effect of flow demand shocks on the real price of oil receded to levels similar to 2007.

In contrast, flow supply shocks and speculative oil demands shock did not play a decisive role during 2003-10. This result is particularly informative given that the same model assigns an important role to speculative demand shocks during 1979, 1986, and 1990/91. These findings indicate that the model indeed has the ability to detect speculation when it exists. Although the structural model cannot distinguish between desirable and undesirable speculation, this distinction does not matter from a policy point of view, given that in the model there is no evidence of speculation being responsible for the recent volatility in the real price of oil. This finding by construction also rules out the hypothesis that undesirable speculation by financial market participants caused these oil price fluctuations.

4. Forecast Scenarios

The reduced-form representation of the VAR model discussed in section 2 may be used to generate real-time out-of-sample forecasts of the real price of oil. The question of interest in this paper is how sensitive these forecasts are to changes in the underlying assumptions about future oil demand and oil supply shocks. This question may be addressed with the help of the structural moving average representation we already used in section 3 to analyze historical fluctuations in the real price of oil.

There is a strict correspondence between standard reduced-form VAR forecasts and forecasts from the structural moving-average representation for a given data set. The reduced-form forecast corresponds to the expected change in the real price of oil conditional on all future shocks being zero in expectation. Departures from this benchmark can be constructed by feeding pre-specified sequences of *future* structural shocks into the structural moving-average representation of the VAR model and projecting the dependent variable into the future. Such sequences of future oil demand and/or oil supply shocks are known as forecast scenarios. Forecast scenarios may be based on sequences of past structural shocks or may be purely hypothetical.

By analogy to equation (1), we can write the structural moving average representation of a vector autoregression as:

$$y_{t+h} = \sum_{i=0}^{\infty} \Theta_i w_{t+h-i} = \sum_{i=0}^{h-1} \Theta_i w_{t+h-i} + \underbrace{\sum_{i=h}^{\infty} \Theta_i w_{t+h-i}}_{y_t}, \quad (2)$$

where y_{t+h} denotes the dependent variable h months in the future. Θ_i and w_t can again be replaced by consistent estimates in practice. Setting all future structural shocks in (2) to zero yields the reduced-form forecast or unconditional VAR forecast. Feeding in a sequence of nonzero future structural shocks provides a conditional forecast. This conditional forecast differs from that discussed in Waggoner and Zha (1999) in that we do not condition on a pre-specified path for observables but rather on a sequence of structural shocks, which is why we refer to this forecast as a forecast scenario.

It is convenient to normalize all conditional forecasts relative to the baseline forecast from the structural model obtained by setting all future structural shocks to zero, which eliminates the dependence of the forecast scenario on y_t . The plot of this normalized conditional forecast represents the upward or downward adjustments of the baseline forecast that would be required if a given hypothetical scenario were to occur.

In other words, for a given sequence of future structural shocks $\{w_{t+1}^{scenario}, \dots, w_{t+h}^{scenario}\}$, the revision required in the baseline forecast of y_{t+h} , $h=1,2,\dots$, if the scenario were to come true, would be:

$$y_{t+h}^{revision} = \sum_{i=0}^{h-1} \Theta_i w_{t+h-i}^{scenario} - \sum_{i=0}^{h-1} \Theta_i w_{t+h-i}^{baseline} = \sum_{i=0}^{h-1} \Theta_i w_{t+h-i}^{scenario}$$

This approach is formally similar to the construction of impulse response functions. The key difference is that an impulse involves a one-time structural shock $w_t^{scenario} \neq 0$, followed by $w_{t+i}^{scenario} = 0 \forall i > 0$, whereas forecast scenarios tend to involve sequences of nonzero structural shocks extending over several periods.

An important difference between conditional and unconditional forecasting methods is that in generating conditional forecasts (like in impulse response analysis) the objective is to identify to the best of our ability the true dynamic effects of a sequence of structural shocks, which necessitates the use of fully revised data. Accordingly, all forecast scenarios in this paper are based on fully revised data,

whereas the unconditional forecasts are based on real-time data. Moreover, given the linearity of the VAR model, forecast scenarios (like impulse response functions) are time invariant. In other words, a given hypothetical sequence of shocks will cause the same revisions of the baseline forecast at each point in time. This also means that forecast scenarios in practice do not necessarily have to be recomputed every month, except to the extent that a longer sample offers efficiency gains in estimating the structural model.

The construction of forecast scenarios allows the forecaster to explore various risks inherent in the forecast of the real price of oil and to convey the consequences to the end-user. In fact, these forecast scenarios could be readily interpreted in conjunction with any baseline forecast, be it a reduced-form VAR forecast or a no-change forecast, for example. In conjunction with historical decompositions, this approach satisfies the needs of central bankers and of other end-users not only for accurate out-of-sample forecasts, but for a coherent economic interpretation of both historical oil price data and oil price forecasts.

5. Examples of Forecast Scenarios

As discussed earlier, forecast scenarios may be based on sequences of structural shocks that occurred during selected episodes in the past or may be purely hypothetical. Below we illustrate both approaches. Forecast scenarios by construction involve sequences of unanticipated shocks. This does not mean that one cannot model the occurrence of anticipated events in the oil market. Such events would be reflected in speculative oil demand shocks in the context of the Kilian and Murphy (2010) model. We will provide two such examples below.

Figure 2 considers five specific scenarios. The forecast horizon in all cases is set to 24 months for illustrative purposes. The first scenario involves an unexpected return of Iraqi oil production to its pre-war level (*Iraq at Full Capacity*). Prior to the invasion of Kuwait in August of 1990, Iraq produced 3.454 millions of barrels per day (mbd). Today, Iraqi production has stabilized at 2.625 mbd. An obvious question of interest to market participants is how much further efforts at increasing Iraqi oil production are likely to relieve upward pressures on the real price of oil. Although there is reason to believe that

neglect may have reduced Iraqi capacity compared with 1990, a useful benchmark is a scenario of an increase in global crude oil production such that Iraq returns to full capacity, which would involve an increase of 0.829 mbd. This corresponds to an increase in global crude oil production of 1.1%, which is well within the variation of historical data. We simulate the effects of such a stimulus by calibrating a one-time structural oil supply shock such that the impact response of global oil production growth in the first forecast period is 1.1%. All other future structural shocks are set to zero. Figure 2 shows that the resulting reduction in the real price of oil expressed in percent relative to the baseline forecast is modest. The real price of oil would temporarily decline by about 5%.

The second scenario involves shutting down Libyan oil production, motivated by events in Libya since February of 2011 (*Libyan Production Shortfall*). This event translates to a negative flow supply shock corresponding to a 2.2% decline in global crude oil production on impact. The real price of oil temporarily increases by as much as 7% relative to the baseline forecast (see Figure 2). This example illustrates that the sustained increase in the real price of oil since February 2011 cannot be attributed to the oil supply disruption alone. Indeed, the model allows for another important channel by which geopolitical events in the Middle East may affect the real price of oil and that channel is speculative demand. As political unrest and civil strife in Egypt, Oman, Yemen, Bahrain, Libya and Syria grows, it is only natural for oil market participants to grow concerned about the political stability of major oil suppliers such as Saudi Arabia, resulting in increased demand for oil inventories and hence a higher real price of oil. As this shift in speculative demand is driven by fears of contagion of political unrest or war, we refer to this situation as a *contagion scenario*. In principle such fears could be arbitrarily weak or strong, making it difficult to assess the quantitative importance of this channel, but the historical experience of earlier episodes in Figure 1 provides some guidance.

One contagion scenario can be motivated by focusing on the surge in speculative demand that occurred preceding and following the invasion of Kuwait in August of 1990 (*Contagion I*). As discussed in Kilian (2008) and Kilian and Murphy (2010), among others, the invasion not only caused oil production in Kuwait and Iraq to cease, but raised concerns that Saudi Arabia and its smaller neighbors

would be invaded next, causing a surge in speculative demand that only subsided after the U.S. had moved troops in sufficient strength to Saudi Arabia to forestall a second invasion, at which point the real price of oil declined sharply. This event provides a blueprint for a dramatic, but temporary surge in speculative demand and in the real price of oil. Figure 2 shows that when feeding in the estimated sequence of speculative demand shocks for March 1990 through March 1991, the real price of oil rises 19% above its baseline forecast, but that dramatic increase is followed by a decline to 11% below the baseline forecast after 15 months, as crude oil inventories are liquidated after the conclusion of the crisis.

An alternative contagion scenario is the possibility of a more sustained speculative frenzy such as that which occurred starting in mid-1979 after the Iranian Revolution, amidst expectations of continued strong global growth (*Contagion 2*). The growing influence of radical Islam and the hostage crisis in Iran led to increasing tension in the Middle East, even after Iranian oil production had resumed. There was a real possibility of a military conflict between Iran and the United States as well as concern that Iran would move against its oil-producing neighbors in the Persian Gulf. The geopolitical risk was further heightened by the Soviet invasion of Afghanistan later in 1979. These events in conjunction with a booming world economy caused sustained increases in inventory demand and in the real price of oil between May and December of 1979 (see Figure 1). Unlike in the Gulf War scenario, there were no indications that the tension would be resolved soon. This second contagion scenario involves feeding into the model future structural shocks corresponding to the sequence of speculative demand shocks that occurred between 1979.1 and 1980.2. Figure 2 shows that an event such as this would raise the baseline forecast by as much as 22% after 15 months. Compared with the earlier contagion scenario the effects would be smaller, but more persistent.

The last panel of Figure 2 considers a scenario that combines the Libyan production shortfall with a contagion scenario built on the 1990/91 contagion scenario. Because structural shocks are mutually uncorrelated, such combined scenarios can be constructed by simply adding the individual scenarios. The model predicts an overall increase in the real price of oil of about 23%. Considering that the refiners' acquisition cost for crude oil imports stood at \$88 in January of 2011, this would have implied an increase

to \$104 by May 2011, with the expectation that the real price might rise as high as \$108 dollars by July if the 1990/91 episode is any guide. Although actual shifts in inventory demand may be higher or lower than this historical precedent, our analysis provides at least a benchmark for thinking about the potential effects of such geopolitical events.

Whereas the scenarios in Figure 2 hypothesize future changes in oil production and/or in speculative demand, an alternative thought experiment involves a recovery of the flow demand for oil and other industrial commodities. Although the historical decomposition in Figure 1 indicates a substantial recovery since the trough of the Great Recession, consistent with renewed robust growth in emerging Asia, this recovery has not been complete. By 2010.6, the cumulative effect of flow demand shocks on the real price of oil was at about the same level as in September 2007. The *global recovery scenario* considered in the upper panel of Figure 3 asks how a further unexpected surge in the demand for oil like the one that occurred during 2007.9-2008.6 would affect the real price of oil. This scenario involves feeding into the structural moving-average representation future flow demand shocks corresponding to the sequence of global flow demand shocks that occurred from 2007.9 to 2008.6, while setting all other future structural shocks equal to their expected value of zero. Figure 3 shows a persistent hump-shaped increase in the real price of oil reaching a peak after a year and a half. The predicted real price of oil exceeds the baseline forecast by up to 50%, underscoring the sensitivity of the real price of oil to global business cycle fluctuations.

Of equal interest is the scenario of a global collapse triggered by the looming Euro crisis. This crisis has been compared to the collapse of Lehman Brothers in September of 2008 (*Lehman*). One way of assessing the effects of such a crisis on the real price of oil is to feed in the sequence of flow demand shocks that occurred between September and December of 2008. Figure 3 shows that a recurrence of such an event would be expected to lower the real price of oil by about 80%, as world demand for oil collapses. This outcome does not seem implausible considering that in 1998, following the Asian crisis, the real price of oil in nominal terms dropped to about \$10 per barrel, and illustrates the great degree of uncertainty surrounding forecasts of the real price of oil.

We conclude with two *nightmare scenarios* which combine the global recovery scenario with the Libyan production shortfall scenario and one of the contagion scenarios. In other words, a *nightmare* refers to a conjunction of scenarios that put upward pressure on the real price of oil. The revisions in the baseline forecast depend primarily on how the contagion is modeled. The first nightmare scenario in the bottom panel of Figure 3 focuses on a gradual and sustained increase in speculative demand, as occurred in 1979. The second nightmare scenario is based on a more abrupt, but shorter-lived episode of speculative demand pressures, as occurred between March 1990 and March 1991. The first scenario implies a dramatic upward adjustment by 58% in the real price of oil after eight months, followed by a decline to levels of about 30%-34% in excess of the baseline forecast after 15 months. To put these results in perspective, suppose that we are generating a forecast as of January 2011 (right before the Libyan crisis), when the price of oil was \$88. Taking the no-change forecast as the benchmark, Figure 3 predicts a surge in the real price of oil to \$174 by August 2011, followed by a decline to between \$143 and \$148 (all in January 2011 prices) by the end of 2012.

The second scenario is associated with a more gradual, but sustained increase in the real price of oil relative to the baseline forecast. The peak increase is 68% above the baseline forecast after about a year, followed by a more gradual decline to 54% above the baseline forecast. To put this scenario in perspective, again consider a no-change forecast of the real price of oil as of January 2011 as the baseline. Figure 3 implies that the real price of oil would be expected to peak at \$185 (in January 2011 prices) in early 2012 before declining to \$169 by the end of 2012. Clearly such forecast scenarios are low probability events, but it is precisely risk assessments along such lines that are needed in policy analysis.

6. Real-Time Forecasts as of December 2010

Based on recursive forecasts starting in 1991.¹² Baumeister and Kilian (2011) showed that the unrestricted VAR(12) reduced-form model including the percent change in global crude oil production, the Kilian index of global real economic activity, the real price of oil and the change in global crude oil inventories, for example, has lower real-time MSPE and significantly higher real-time directional

accuracy than alternative forecasting models including the no-change forecast, a forecast based on oil futures prices, and univariate AR and ARMA model forecasts. The real-time MSPE reductions relative to the random walk model may be as high as 25% at the one-month horizon and as high as 19% at the three-month horizon. Gains at longer horizons tend to be smaller. The gains in directional accuracy may be as high as 10% even at horizons of one year. For the analysis in this paper, we nevertheless rely on the VAR(24) real-time forecast because of the longer estimation sample available as of 2010.12 compared with the recursive estimation windows in Baumeister and Kilian (2011). The VAR(24) real-time forecast shown in Figure 4 is qualitatively similar to the VAR(12) real-time forecast based on the same data, making that distinction largely moot.

The use of real-time data in forecasting raises two distinct complications. First, even preliminary data may become available only with a lag, necessitating the use of nowcasting techniques in filling gaps in the preliminary data. Second, past data are continuously revised for several months. In Figure 4, the vertical lines indicate the end of the nowcasting period. Based on the preliminary data available to the forecaster in 2010.12, the real price of oil in 2010.12 was \$97. The actual price was closer to \$90. Taking account of these real-time data limitations (and ignoring the effects of the unpredictable Libyan crisis), the VAR model forecast in Figure 4a appears remarkably accurate compared with the no-change forecast and the forecast based on oil futures prices. The latter forecast is constructed by extrapolating the current real price of oil based on the nominal oil futures spread adjusted for expected inflation (see Baumeister and Kilian 2011).

The VAR model forecasts an initial increase in the real price of oil to \$105 after one quarter even in the absence of the Libyan crisis, followed by decline to between \$77 and \$83 in the second year. This matches the general evolution of the real price of oil since 2010.12 to date. The U.S. refiners' acquisition cost for crude oil imports peaked in April at \$113. The even higher increase in the real price of oil presumably reflects the Libyan oil supply crisis that began in late February of 2011. The EIA to date has not yet issued the preliminary data for the refiners' acquisition cost for crude oil imports beyond July, but the corresponding nominal WTI price as of September is \$86, consistent with the prediction of a gradual

price decline. The futures-based forecast also indicates a gradual decline after two years, but only from \$97 to \$90. That decline is largely driven by expected inflation. The term structure of the nominal oil futures price is nearly flat across horizons.

Figure 4b shows the full range of alternative forecast scenarios discussed in section 5. As in Figure 4a, all forecasts in Figure 4b are expressed in 2010.12 dollars. For expository purposes, we postulate that all scenarios start in January of 2011. One could relax that assumption and delay the onset of a given scenario if appropriate. In this section, the focus is on illustrating how sensitive the baseline forecast is to alternative assumptions about future demand and supply conditions in the global market for crude oil. Figure 4b illustrates that the real price of oil may rise as high as \$148 after one quarter or fall as low as \$66. After one year, the range is between \$19 and \$133; at the two-year horizon between \$26 and \$120. These results illustrate how structural models of oil markets may be used to assess the sensitivity of reduced-form forecasts of the real price of oil to alternative assumptions about future oil demand and oil supply shocks. The next section discusses how this analysis may be combined with a formal analysis of the forecast risks to generate additional insights.

7. Real-Time Risk Analysis

Alternative forecast scenarios can be combined into one conditional forecast, if one is willing to attach probability weights to each individual scenario. A useful tool in constructing the probability weights is a Venn diagram that illustrates the relationship between alternative scenarios. An example of such a diagram is shown in Figure 5. There are six distinct forecast scenarios: *baseline*, *Lehman*, *global recovery*, *Libya*, *contagion 1* and *contagion 2*. The *Iraq at full capacity* scenario has been omitted to simplify the presentation of the results. Effectively, we treat this scenario as having zero probability.

We impose that the Lehman and global recovery scenarios are mutually exclusive. We also treat the baseline scenario as mutually exclusive with the other scenarios. Finally, we postulate that the two contagion scenarios are mutually exclusive. Subject to these restrictions, the diagram allows for any given scenario to occur in conjunction with one or more of the other scenarios. Having made explicit the

relationship between the individual scenarios, the end-user must specify the probabilities associated with each event in the diagram. These probabilities must sum to one.

Table 2 provides four examples. In the baseline forecast, the *baseline scenario* receives weight 1 and the other scenarios zero weight. Alternative 1 refers to a moderately pessimistic scenario, in which the end-user attaches considerable probability weight to adverse events in the oil market and a relatively low probability weight to a global recovery. Alternative 2 represents a more pessimistic scenario in which the probability weight attached to the *Lehman scenario* is increased at the expense of the *baseline scenario*. Finally, alternative 3 is an example of a forecaster who is more optimistic about the economy, reflected in a higher probability weight on the global recovery scenario and lower probability weights on adverse events in the oil market.

Given the fitted VAR model, it is straightforward to compute the real-time predictive density for the baseline forecast by bootstrap simulation methods. The use of such methods for direct forecasts is discussed in Alquist, Kilian, and Vigfusson (2011). For iterated forecasts from VAR models the standard approach suffices. In resampling the residuals, we rely on bootstrap methods. In practice, we use 20,000 random draws to simulate the predictive distribution for a given horizon conditional on the estimated parameters.

The construction of predictive densities for alternative scenarios is facilitated by the fact that departures from the baseline scenario do not change the shape of the predictive distribution. They merely shift the real-time predictive density by some constant, given by the dollar amount by which the baseline forecast differs from the forecast path under the alternative scenario, as illustrated in Figure 4b. This is a consequence of the linearity of the VAR model.

The probability weights shown in Table 1 allow us to assign a probability to each area in the Venn diagram. For any combination of events in Figure 5, we can compute the implied predictive density for the real price of oil in 2010.12 dollars, by simply adding the appropriate dollar amount to the baseline predictive density. The predictive density of the weighted scenario then corresponds to the probability-weighted densities for each area in the Venn diagram.

Figure 6 displays the one-year ahead real-time predictive density for the baseline scenario as well as the corresponding densities for the three weighted alternative scenarios described in Table 1. Figure 6 illustrates the importance of the choice of probability weights. Alternative views about the likelihood of alternative scenarios can result in substantial shifts in probability mass. In assessing the implications of these weighted forecast densities, it is useful to conduct a formal analysis of the forecast risks. A risk measure by definition depends both on the predictive distribution of the variable of interest and on the preferences (or loss function) of the end-user of the forecast (see Machina and Rothschild 1987). Our risk analysis adapts the tools discussed in Alquist, Kilian, and Vigfusson (2011). Let R_{t+h} denote the real price of oil in 2010.12 dollars h periods from date t and $F(\cdot)$ the predictive distribution of R_{t+h} . Consider the events of R_{t+h} exceeding an upper threshold of \bar{R} (upside risk) and of R_{t+h} falling below the lower threshold of \underline{R} (downside risk). For example, we may consider $\underline{R} = 80$ and $\bar{R} = 100$. Let α and β denote the forecast user's degree of risk aversion. In that case, under weak assumptions on the forecast user's preferences,

$$DR_{\alpha} \equiv -\int_{-\infty}^{\underline{R}} (\underline{R} - R_{t+h})^{\alpha} dF(R_{t+h}), \quad \alpha \geq 0$$

$$UR_{\beta} \equiv \int_{\bar{R}}^{\infty} (R_{t+h} - \bar{R})^{\beta} dF(R_{t+h}), \quad \beta \geq 0$$

represent the downside risk and upside risk embodied in the forecast, respectively.¹

Under risk neutrality $\alpha = \beta = 0$, the risk measures above reduce to the *target probabilities* $DR_0 = -\Pr(R_{t+h} < \underline{R})$ and $UR_0 = \Pr(R_{t+h} > \bar{R})$. For $\alpha = \beta = 1$, the risk measures above can be shown to reduce to the *tail conditional expectations* $DR_1 = E(R_{t+h} - \underline{R} | R_{t+h} < \underline{R})\Pr(R_{t+h} < \underline{R})$ and $UR_1 = E(R_{t+h} - \bar{R} | R_{t+h} > \bar{R})\Pr(R_{t+h} > \bar{R})$. The tail conditional expectation is not only concerned with the likelihood of a tail event, but also with how far the real price of oil is expected to be in the tail if that tail event occurs. The latter expression is also known as the *expected shortfall* (or expected excess).

¹ Measures of risk of this type were first introduced by Fishburn (1977), Holthausen (1981), Artzner, Delbaen, Eber and Heath (1999), and Basak and Shapiro (2001) in the context of portfolio risk management and have become a standard tool in recent years. For additional discussion of the risk measures above see Kilian and Manganelli (2007, 2008).

Extensions to higher degrees of risk aversion would be straightforward.

Below we report the tail conditional expectation as well as its two components based on predictive densities generated from real-time data for 2010.12. Analogous risk measures have been employed by policymakers and financial analysts in other contexts (see Kilian and Manganeli 2007). In constructing these risk measures from the predictive distribution of the real price of oil, the expectations and probabilities in question can be estimated by their sample equivalent. Table 2 shows the tail conditional expectation and its components for the four predictive densities displayed in Figure 6.

It is useful to start with the baseline scenario. The probability of the real price of oil exceeding \$100 is 67% at the 3-month horizon, but declines to 46% at the 6-month horizon and to about 20% at longer horizons. This pattern is mirrored by an increase in the probability of the real price of oil dropping below \$80. At the 3-month horizon, that probability is only 2%, but at horizons of one year or longer it rises to about 50%. Probabilities alone however tell us nothing about how far the real price of oil is expected to be outside the forecast user's comfort zone somewhere between \$80 and \$100. This question is answered by the expected shortfall and expected excess measures shown in the next column. For example, the expected excess price relative to the \$100 benchmark is about \$14 at the 3-month horizon, but rises to \$23 at the two-year horizon. Unlike the probability of exceeding \$100, the corresponding expected excess measure is monotonic in the horizon. A risk measure that aggregates both of these elements of risk is the tail conditional expectation shown in the next column. Table 6 indicates that overall the upside oil price risk declines at horizons up to one year with a partial recovery at the two-year horizon, whereas downside risks monotonically increase with the horizon. At the one-year horizon, for example, downside risks outweigh upside risks by a factor of about 2:1.

Whereas the alternative moderately pessimistic scenario 1 is quite similar to the baseline scenario, scenario 2 embodies a much more pessimistic outlook and hence indicates a somewhat lower upside risk. At the 3-month horizon, the probability of the real price of oil exceeding \$100 is only 61%; at the 6-month horizon it drops to 26%. The corresponding probabilities of the real price of oil falling below \$80 rise from 3% at the 3-month horizon to 36% at the 6-month horizon. Taking account of the expected

shortfall and expected excess, beyond the 3-month horizon, overall downside risks outweigh overall upside risks. The risk imbalances are most striking at longer horizons. The optimistic scenario 3, on the other hand, illustrates that a substantial increase in the probability of an unexpected global recovery could overcome the implication of the baseline model that the real price of oil is expected to decline in 2012. In this case, at all horizons the odds of the real price of oil exceeding \$100 exceed those of the real price dropping below \$80. The tail conditional expectation favors the upside risk at the two-year horizon with a factor of 5:1, also reflecting the substantially larger expected excess on the upside.

The tools presented in this section are designed to make explicit the tradeoffs between alternative assumptions. They also make explicit by how much and at what horizon downside and upside risks change as a function of the probability weights assigned to each individual scenario. Clearly, the probability of a global recovery on the one hand and that of a Lehman scenario (say triggered by a European banking crisis) on the other, are key parameters in assessing the risks embodied in forecasts of the real price of oil, as is the probability of contagion scenarios. Scenarios involving oil production alone, on the other hand, are less important.

8. Conclusion

Users of forecasts of the real price of oil such as central banks are interested not only in accurate out-of-sample forecasts, but also in understanding the past, current, and future evolution of the real price of oil. Of particular importance is the ability to quantify the risks associated with a baseline forecast based on an extensive analysis of how this forecast changes under alternative hypothetical forecast scenarios. We illustrated how both of these objectives may be attained using an oil market VAR model that recently has been shown to forecast the real price of oil more accurately than alternative forecasting models. With suitable additional identifying assumptions, this VAR model may be used not only to interpret past and current fluctuations in the real oil price data in light of economic models, but to evaluate the sensitivity of the baseline forecast to alternative forecast scenarios. Such scenarios involve alternative assumptions about sequences of future oil demand and oil supply shocks. We showed how forecast scenarios may be

generated from the structural moving-average representation of the VAR model and combined with real-time forecasts of the real price of oil.

For expository purposes, we explored six basic scenarios involving hypothetical future oil demand and oil supply conditions as well as scenarios for which there is historical precedent. We showed that an unexpected recovery of the world economy (comparable to the surge in global flow demand for oil from mid-2007 to mid-2008), for example, would raise the real price of oil by an additional 50% after a year and a half. On the other hand, a surge in speculative demand driven by civil unrest in the Middle East comparable to that during the Iranian crisis of 1979 would increase the real price of oil by 20% after about one year.

In order to fully appreciate the implications of these and other forecast scenarios, we proposed that end-users of oil price forecasts construct probability-weighted real-time forecast densities for the real price of oil. We showed how to explore the implications of changes in the probability weights attached to various forecast scenarios and combinations of forecast scenarios using formal risk measures, building on the analysis in Kilian and Manganelli (2007, 2008). Such risk measures may be computed from the predictive density of the real price of oil. Risk measures facilitate the interpretation of the information conveyed by the predictive density. We illustrated how to identify the key parameters driving upside and downside risks in the forecast of the real price of oil.

This analysis can only be conducted within the context of a specific structural model. The structural oil market model of Kilian and Murphy (2010) we used in this paper allowed us to contrast the role of flow supply shocks, flow demand shocks and speculative demand shocks in oil markets, but is not without limitations. It does not, for example, allow a further distinction between risks arising from different monetary or fiscal policies or between risks associated with unexpected changes in oil intensity and in aggregate levels of production. More elaborate risk analysis along these lines would require a dynamic stochastic general equilibrium (DSGE) model. An example of such a model would be Bodenstein and Guerrieri (2011). Of course, these benefits come at the expense of requiring much more explicit assumptions about the structure of the global economy and of the market for crude oil. The


guiding principles underlying our analysis of the risks embodied in real oil price forecasts, however, remain the same regardless of the choice of the structural model.

References

- Alquist, R., and Kilian, L. (2010), “What Do We Learn from the Price of Crude Oil Futures?” *Journal of Applied Econometrics*, 25, 539-573.
- Alquist, R., Kilian, L., and Vigfusson, R.J. (2011), “Forecasting the Price of Oil,” prepared for: G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting*, 2, Amsterdam: North-Holland.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999), “Coherent Measures of Risk,” *Mathematical Finance*, 9, 203-228.
- Basak, S., and Shapiro, A. (2001), “Value-at-Risk Based Management: Optimal Policies and Asset Prices,” *Review of Financial Studies*, 14, 371–405.
- Baumeister, C., and Kilian, L. (2011), “Real-Time Forecasts of the Real Price of Oil,” forthcoming: *Journal of Business and Economic Statistics*.
- Baumeister, C., and Peersman, G. (2011), “The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market,” mimeo, Department of Economics, Ghent University.
- Bodenstein, M., and Guerrieri, L. (2011), “Oil Efficiency, Demand, and Prices: A Tale of Ups and Downs,” mimeo, Federal Reserve Board.
- Fishburn, P.C. (1977), “Mean-Risk Analysis with Risk Associated with Below-Target Returns,” *American Economic Review*, 67, 116–26.
- Giannone, D., and Reichlin, L. (2006), “Does Information Help Recover Structural Shocks from Past Observations,” *Journal of the European Economic Association*, 4, 455-465.
- Hamilton, J.D. (2009), “Understanding Crude Oil Prices”, *Energy Journal*, 30, 179-206.


- Holthausen, D.M. (1981), “A Risk-Return Model with Risk and Return Measured in Deviations from Target Return,” *American Economic Review*, 71, 182–88.
- Kellogg, R. (2011), “Short-Run Responses of Oil Production to Spot and Future Oil Prices,” mimeo, University of Michigan.
- Kilian, L. (2008), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” *Review of Economics and Statistics*, 90, 216-240.
- Kilian, L. (2009), “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., and Manganelli, S. (2007), “Quantifying the Risk of Deflation,” *Journal of Money, Credit, and Banking*, 39, 561-590.
- Kilian, L., and Manganelli, S. (2008), “The Central Banker as a Risk Manager: Estimating the Federal Reserve’s Preferences under Greenspan,” *Journal of Money, Credit and Banking*, 40, 1103-1129.
- Kilian, L., and Murphy, D.P. (2010), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” mimeo, University of Michigan.
- Kilian, L., and Murphy, D.P. (2011), “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models,” forthcoming: *Journal of the European Economic Association*.
- Lütkepohl, H. (2005), *New Introduction to Multiple Time Series Analysis*, Springer: New York.
- Machina, M.J., and M. Rothschild (1987), “Risk,” in Eatwell, J., Millgate, M., and P. Newman (eds.), *The New Palgrave Dictionary of Economics*, London: MacMillan, 203-205.
- Waggoner, D.F., and Zha, T. (1999), “Conditional Forecasts in Dynamic Multivariate Models,” *Review of Economics and Statistics*, 81, 639-651.
- Yergin, D. (1992), *The Prize. The Epic Quest for Oil, Money, and Power*, New York: Simon and Schuster.

Figure 1: Historical Decomposition for Real Refiners' Acquisition Cost from Structural VAR(24) Model



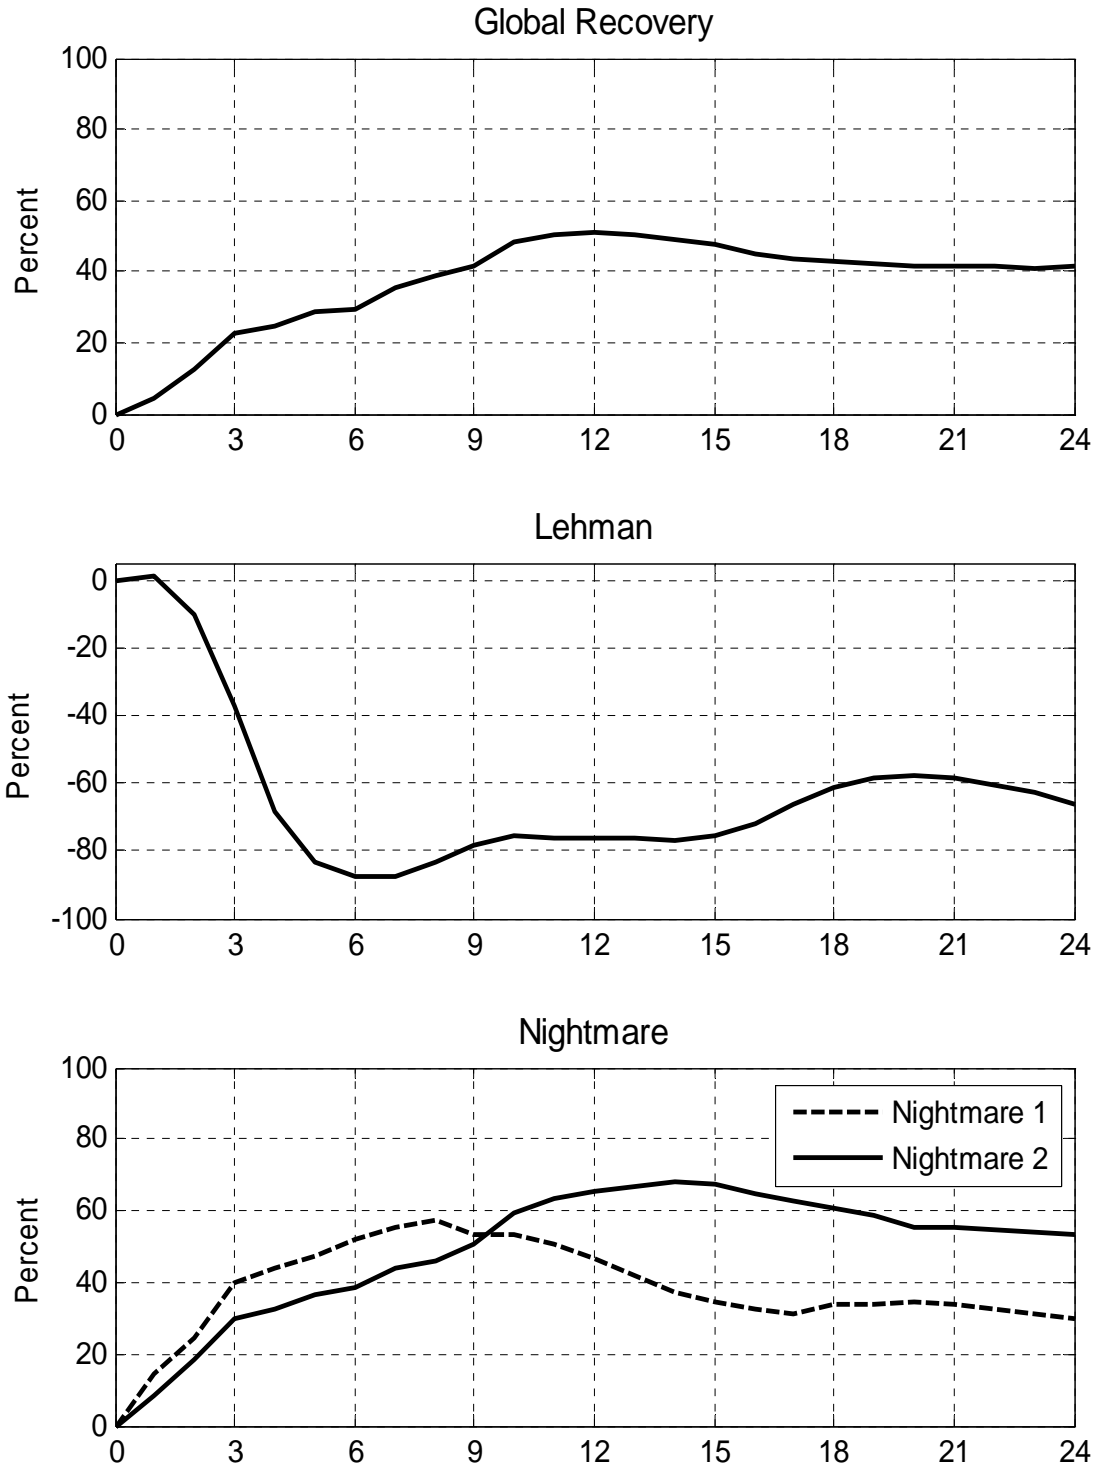
NOTES: The vertical bars correspond to key events in the global market for crude oil including the Iranian Revolution in late 1978, the outbreak of the Iran-Iraq War in late 1980, the collapse of OPEC at the end of 1985, the August 1990 invasion of Kuwait, the Asian crisis of 1997, the Venezuelan oil supply shock of late 2002 (immediately followed by the Iraq War of early 2003), and the global financial crisis of mid-2008.

**Figure 2: Forecast Scenarios for Real Refiners' Acquisition Cost
Percent Deviations from Baseline Forecast**




NOTES: A description of each scenario can be found in section 5.

**Figure 3: Forecast Scenarios for Real Refiners' Acquisition Cost
Percent Deviations from Baseline Forecast**




NOTES: The two *nightmare scenarios* combine the global recovery scenario with the Libyan production shortfall scenario and with the contagion 1 and contagion 2 scenarios, respectively.

Figure 4: Baseline Real-Time Forecast of Real Refiners' Acquisition Cost as of 2010.12 and Risk Analysis




NOTES: The scenarios in panel (b) correspond to those in Figures 3 and 4. The vertical line indicates the last nowcast.

**Figure 5. Event Analysis using Venn Diagrams
An Illustrative Example**



NOTES: B stands for *baseline*, Le for *Lehman*, Li for *Libyan production shortfall*, R for *global recovery*, C₁ and C₂ stand for *contagion 1* and *contagion 2*. We abstract from the *Iraq at full capacity* scenario for expository purposes. The Le and R scenarios are mutually exclusive, as is the baseline scenario with the other scenarios. Likewise C₁ and C₂ are treated as mutually exclusive.

**Figure 6. Real-Time Probability-Weighted 1-Year Ahead Predictive Densities for the Real Price of Oil as of 2010.12
An Illustrative Example**



NOTES: The probability weights for the baseline scenario and for the alternative scenarios are shown in Table 1.

Table 1: Probability Weights for Forecast Scenarios

Events	Baseline	Weighted Forecast Scenarios		
		Alternative 1: Moderately Pessimistic	Alternative 2: Pessimistic on Economy	Alternative 3: Optimistic on Economy
B	1	0.41	0.31	0.16
C ₂	0	0.16	0.16	0.11
Li∩C ₂	0	0.05	0.05	0.05
Li∩C ₂ ∩R	0	0.03	0.03	0.03
C ₂ ∩R	0	0.04	0.04	0.04
Li∩C ₁ ∩R	0	0.03	0.03	0.03
Li∩C ₁	0	0.05	0.05	0.05
C ₁	0	0.16	0.16	0.11
R∩C ₁	0	0.04	0.04	0.04
Li∩R	0	0.07	0.07	0.07
Le	0	0.13	0.23	0.08
Le∩Li	0	0.03	0.03	0.03
Le∩C ₁	0	0.01	0.01	0.01
Le∩C ₁ ∩Li	0	0.01	0.01	0.01
Le∩C ₂	0	0.01	0.01	0.01
Le∩C ₂ ∩Li	0	0.01	0.01	0.01
R	0	0.14	0.14	0.59
Li	0	0.22	0.22	0.17
I	0	0	0	0

NOTES: B stands for *baseline*, Le for *Lehman*, Li for *Libyan production shortfall*, R for *global recovery*, I stands for *Iraq at full capacity*, C₁ and C₂ stand for *contagion 1* and *contagion 2*. The Le and R scenarios are mutually exclusive, as is the baseline scenario with the other scenarios. Likewise C₁ and C₂ are treated as mutually exclusive. The probabilities have been chosen for expository purposes. For a visual representation see Figure 5. The total probability mass for each weighted scenario is one by construction.

Table 2: Risk Measures for Probability Weighted Forecast Scenarios

Scenario	h	$P(R_{t+h} > 100)$	$E(R_{t+h} - 100 R_{t+h} > 100)$	$E(R_{t+h} - 100 R_{t+h} > 100) \times \Pr(R_{t+h} > 100)$	$P(R_{t+h} < 80)$	$E(R_{t+h} - 80 R_{t+h} < 80)$	$E(R_{t+h} - 80 R_{t+h} < 80) \times \Pr(R_{t+h} < 80)$
Baseline	3	0.67	13.53	9.06	0.02	5.34	0.11
	6	0.46	17.09	7.93	0.15	8.22	1.25
	12	0.20	16.70	3.27	0.51	14.37	7.26
	24	0.23	23.13	5.37	0.52	18.32	9.59
1	3	0.72	14.25	10.19	0.02	5.11	0.08
	6	0.41	16.63	6.79	0.19	8.82	1.71
	12	0.18	16.50	3.02	0.53	14.97	7.94
	24	0.21	22.93	4.82	0.56	19.53	10.94
2	3	0.61	12.72	7.72	0.03	5.29	0.16
	6	0.26	15.10	3.89	0.36	11.50	4.09
	12	0.12	15.81	1.88	0.66	18.86	12.50
	24	0.16	22.40	3.60	0.65	23.07	14.95
3	3	0.93	21.01	19.54	0.00	3.15	0.01
	6	0.74	21.82	16.15	0.03	6.04	0.18
	12	0.60	21.16	12.64	0.08	6.41	0.51
	24	0.47	25.56	11.92	0.23	10.35	2.35

NOTES: The predictive densities from the iterated model were simulated using bootstrap methods.