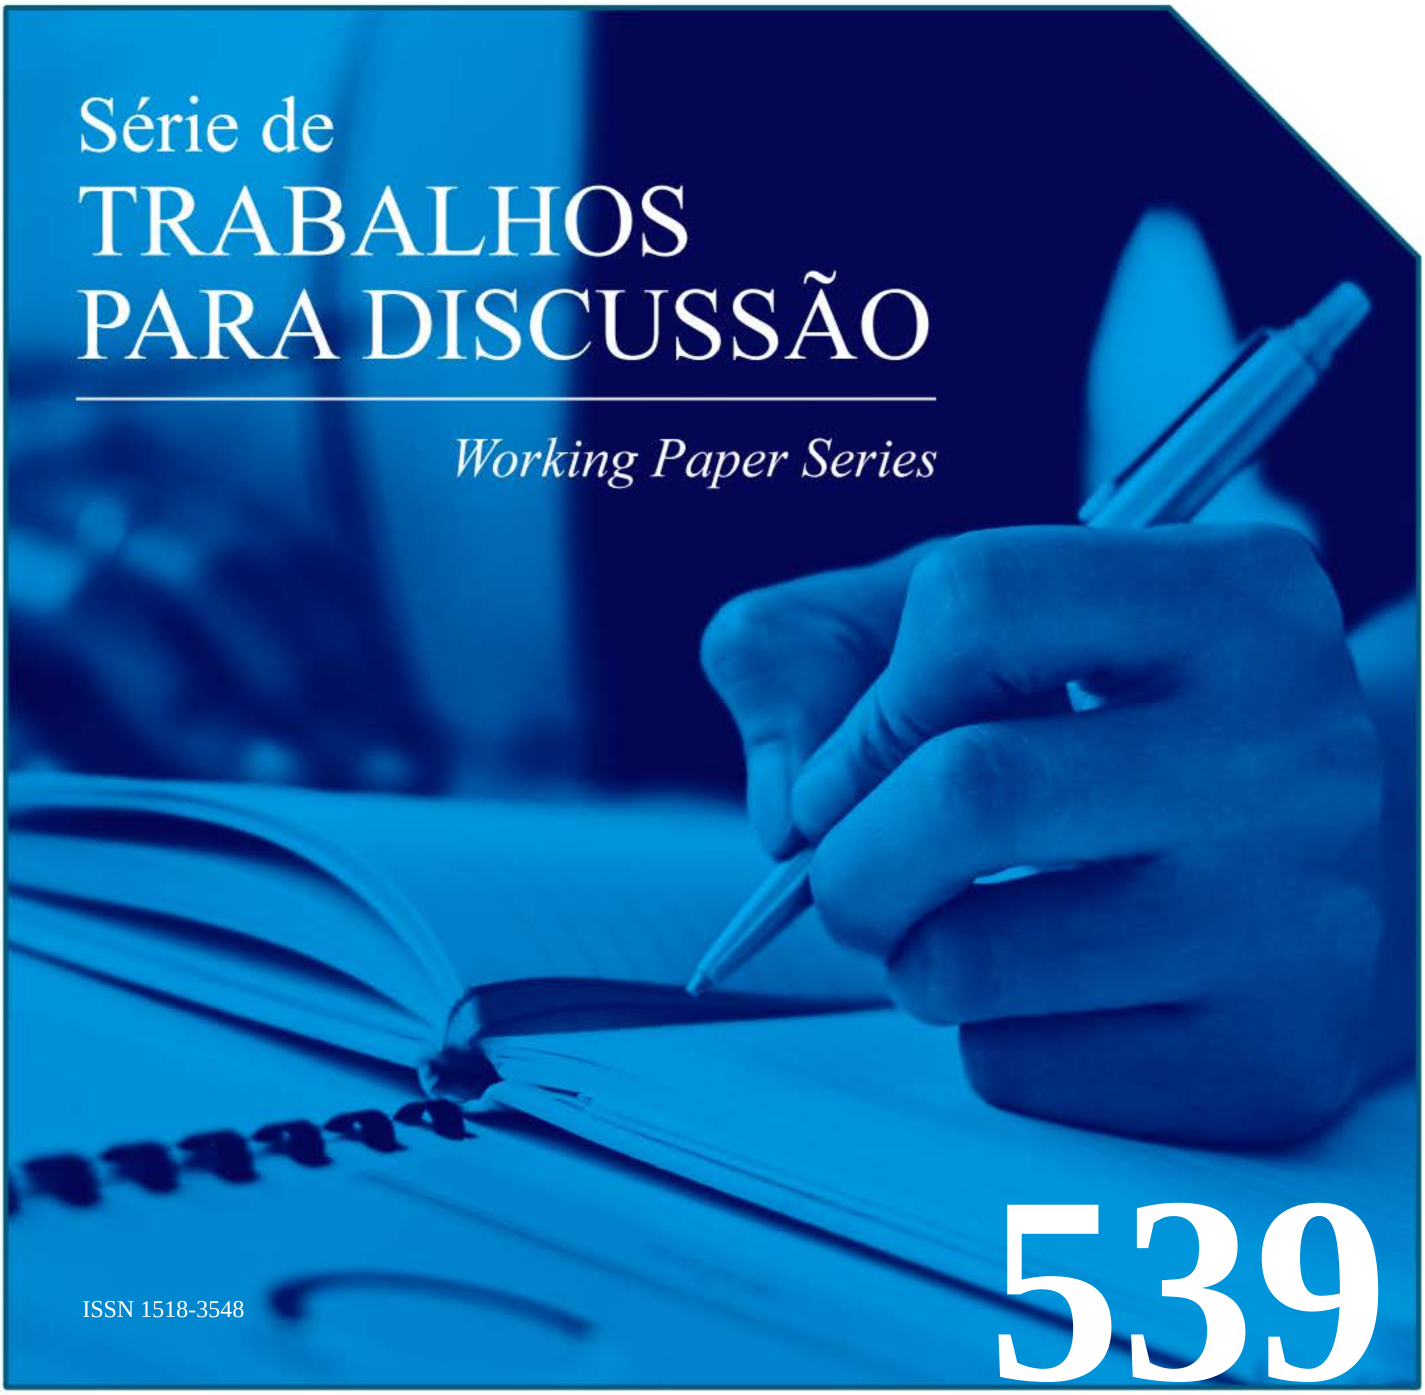

NOVEMBER 2020

Commodity Prices and Global Economic Activity: a derived-demand approach

Angelo Mont'Alverne Duarte, Wagner Piazza Gaglianone, Osmani Teixeira de Carvalho Guillén, João Victor Issler

BANCO CENTRAL DO BRASIL

ISSN 1518-3548
CGC 00.038.166/0001-05

Working Paper Series	Brasília	no. 539	November	2020	p. 1-44

Working Paper Series

Edited by the Research Department (Depep) – E-mail: workingpaper@bcb.gov.br

Editor: Francisco Marcos Rodrigues Figueiredo

Co-editor: José Valentim Machado Vicente

Head of the Research Department: André Minella

Deputy Governor for Economic Policy: Fabio Kanczuk

Citizen Service Division

Banco Central do Brasil

Deati/Diate

SBS – Quadra 3 – Bloco B – Edifício-Sede – 2º subsolo

70074-900 Brasília – DF – Brazil

Toll Free: 0800 9792345

Fax: +55 (61) 3414-2553

Internet: http//www.bcb.gov.br/?CONTACTUS

Non-technical Summary

This paper studies the interaction between commodity prices and global economic activity in a setup where a representative cost-minimizing firm chooses optimal inputs as derived demands. Our focus is on important globally traded commodities, whose supply function is very price inelastic in the short run. Key examples studied here are Oil and major metal commodities, such as Aluminum, Copper and Nickel. Our empirical evidence fully supports the theoretical results of the derived-demand model. Indeed, this paper shows overwhelming evidence that cycles in oil prices are synchronized to those of global industrial production. This evidence is stronger regarding the global economy but holds as well for the U.S. economy. Our first original contribution is to investigate and find common cycles accounting for theory and empirics.

Our second contribution is to investigate the usefulness of the common-feature VAR-based approach for forecasting global measures of economic activity - U.S. GDP and OECD's GDP. In this sense, commodity prices are employed to help forecasting economic activity, relying again on the role commodities play in the derived-demand model. This is important, since commodity prices are observed on an almost continuous-time basis, but economic activity is observed at a much lower frequency. This way, commodity prices could give early signals about future global economic activity. Our empirical evidence in forecasting global measures of economic activity emphasizes the importance of conditioning on either oil prices or metal-commodity prices, relying on parameter-reduction techniques that favor optimal forecasts. Forecasting gains are substantial for U.S. GDP, but results for OECD's GDP are disappointing.

Finally, the third objective of the paper is to forecast oil prices over short horizons (monthly data) based on techniques that view the optimal forecast as a common feature, which can be identified by using a cross-sectional average of a group of forecasts. The model combination approach considers 254 macroeconomic and financial covariates available from different databases. From this group, 20 covariates are selected and used on 32,412 different models - all predicting oil prices. When these models are combined using optimal techniques, they outperform the random-walk model with an out-of-sample R^2 statistic that can reach up to 14%. This is a major improvement vis-a-vis the previous literature.

Sumário Não Técnico

Este artigo estuda a interação entre os preços de *commodities* e atividade econômica global em um arcabouço onde uma firma representativa que minimiza custos escolhe seus insumos de forma ótima. Nosso foco está em *commodities* comercializadas globalmente, cuja função de oferta é muito inelástica ao preço no curto prazo. Os principais exemplos estudados aqui são petróleo e as principais *commodities* metálicas, como alumínio, cobre e níquel. Nossa evidência empírica apoia inteiramente os resultados teóricos do modelo de demanda derivada. De fato, este artigo apresenta evidências contundentes de que os ciclos dos preços do petróleo são sincronizados com os da produção industrial global. Essa evidência é mais forte em relação à economia global, mas também é válida para a economia dos EUA. A primeira contribuição original deste artigo é investigar e encontrar ciclos comuns sob os aspectos teórico e empírico.

Nossa segunda contribuição é investigar a utilidade da abordagem de *common features* baseada em VAR na previsão de medidas globais de atividade econômica - PIB dos EUA e PIB da OCDE. Neste sentido, preços de *commodities* são utilizados para ajudar a prever a atividade econômica, contando novamente com o papel que as *commodities* desempenham no modelo de demanda derivada. Isso é importante, uma vez que os preços das *commodities* são observados em uma base quase contínua, enquanto que a atividade econômica é observada com uma frequência muito menor. Dessa forma, os preços das *commodities* podem dar sinais antecedentes sobre a atividade econômica global futura. Nossa evidência empírica na previsão de medidas globais de atividade econômica enfatiza a importância de condicionar os preços do petróleo ou das *commodities* metálicas, contando com técnicas de redução de parâmetros que favoreçam previsões ótimas. Os ganhos de previsão são substanciais para o PIB dos EUA, mas os resultados para o PIB da OCDE são decepcionantes.

Finalmente, o terceiro objetivo do artigo é prever os preços do petróleo em horizontes curtos (dados mensais) com base em técnicas que consideram a previsão ótima como uma *common feature*, que pode ser identificada por meio de uma média *cross-sectional* de um grupo de previsões. Tal abordagem de combinação de modelos considera 254 covariáveis macroeconômicas e financeiras de diferentes bases de dados. Desse grupo, 20 covariáveis são selecionadas e usadas em 32.412 modelos diferentes - todas prevendo os preços do petróleo. Quando esses modelos são combinados usando técnicas de otimização, eles superam o modelo de passeio aleatório com uma estatística R^2 fora da amostra que pode chegar a 14%, representando uma grande melhoria em relação à literatura anterior.

Commodity Prices and Global Economic Activity: A Derived-Demand Approach*

Angelo Mont'Alverne Duarte†
Wagner Piazza Gaglianone‡
Osmani Teixeira de Carvalho Guillén§
João Victor Issler¶

Abstract

In this paper, a derived-demand approach is proposed to explain the positive correlation and the synchronicity between the growth rates of commodity prices and of economic activity at the global level. The focus is on important traded commodities, whose supply function is very price inelastic in the short run, such as oil and major metal commodities. The paper contributions are as follows. First, the synchronicity of oil-price and global activity cycles is presented using the tools of the common-feature literature. Second, it is shown how to improve forecasts of global activity using commodity prices, noting that one observes the latter at an almost continuous-time basis, but the former at a much lower frequency and with considerable delay. Third, the usefulness of optimal forecast combinations for oil prices is discussed employing a wide array of macroeconomic and financial variables. The out-of-sample R^2 statistic for model combinations can reach up to about 14%, a major improvement over the previous literature.

Keywords: Commodity prices; Derived-demand model; Common features; Oil price forecasts.

JEL Classification: C30; C53; E27; E37.

---

*The views expressed in the paper are those of the authors alone and do not necessarily reflect those of the Banco Central do Brasil or of FGV. We gratefully acknowledge the financial support of CNPq, FGV, FAPERJ, CAPES, and INCT. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior – Brasil (CAPES) – Finance Code 001. We thank Marcia Waleria Machado and Andrea Machado for excellent research assistance.

†Department of Competition and Financial Market Structure, Banco Central do Brasil. E-mail: angelo.duarte@bcb.gov.br

‡Research Department, Banco Central do Brasil. E-mail: wagner.gaglianone@bcb.gov.br

§Open Market Operations Department, Banco Central do Brasil, and IBMEC/RJ. E-mail: osmani.guillen@bcb.gov.br

¶Corresponding Author, Graduate School of Economics and Finance, Getulio Vargas Foundation – FGV EPGE. E-mail: Joao.Issler@fgv.br

1 Introduction

The purpose of this paper is to study the interaction between commodity-price changes and global economic activity measures in a setup where a representative cost-minimizing firm chooses optimal inputs (commodity quantum) as derived demands. The focus is on important globally traded commodities, whose supply function is very price inelastic in the short run. Key examples studied in this paper are Oil (Oil derivatives) and also major metal commodities, such as Aluminum, Copper, and Nickel. The reason for an inelastic short-run supply function lies on the fact that the production function of these commodities is very capital intensive and it takes *time to build capital* (Kydland and Prescott, 1982). So, as a first approximation, the short-run supply of these commodities, which includes inventories, can be treated as given.[1] As production of the global representative firm increases (decreases), there is a contemporaneous increase (decrease) in the prices of these commodities, since their respective supply does not respond to such demand increases (decreases). This leads to potential synchronization between the cycles of commodity prices and those of global economic activity – i.e., *common cycles*.

The econometric tools used in this paper come from the *common feature* literature, e.g., Vahid and Engle (1993), Engle and Issler (1995), Issler and Vahid (2001), Hecq et al. (2006), Issler and Lima (2009), Athanasopoulos et al. (2011), and Gaglianone and Issler (2019), *inter alia*. Unobserved common components of a group of economic series are identified as functions of observables, allowing the consistent estimation of common trends, common cycles, and of optimal forecasts of economic data. In a traditional time-series context, trends and cycles can be jointly investigated in a unified multivariate setting based on vector autoregressive (VAR) models, that could be estimated either by maximum likelihood or by the generalized method of moments. Optimal forecasts, on the other hand, are obtained in panel-data framework, where forecasts of an increasing number of models or of forecast-survey results are optimally combined as either simple averages or weighted averages, respectively. A key property of all of these common unobserved components is that they can be removed by linear combination of the underlying observable series.

The first original contribution of this paper is to test for the synchronization (common cycles) of oil-price and global economic activity variation by using the VAR-based framework. The importance of demand factors for commodity prices has been known at least since the influential work of Deaton and Laroque (1996). For oil prices, recent work on demand factors include Dvir and Rogoff (2014), Kilian and Lee (2014), and Aastveit, Bjørnland and Thorsrud (2015). For metal commodities, Jerrett and Cuddington (2008) conjectured whether there could be super-cycles in their prices linked to the impressive economic growth in China since the 1990's. Issler et al. (2014), on

[1] An inelastic short-run supply function is not without controversy. See Kilian (2009), Kilian and Murphy (2012), Baumeister and Hamilton (2019) and Herrera and Rangaraju (2020).

the other hand, proposed the use of the *derived-demand* model to explain the synchronization between metal-commodity prices and global economic activity. To the best of our knowledge, despite the existence of a mature literature for oil prices, this is the first paper to use modern econometric techniques designed to detect common cycles to examine the synchronicity between the cycles in oil prices and global business cycles. Since oil is *the* major commodity traded in global markets, it is important to understand and to properly model its dynamic behavior using modern econometric tools.

The VAR-based econometric setup used in this paper potentially implies a substantial parameter reduction in these models, leading to efficiency in estimation of the endogenous series of the system. From an empirical point-of-view, the application of these techniques has proven to be useful in different contexts: Issler and Vahid (2001) find a 25% reduction in out-of-sample mean-squared forecast error (MSFE) for U.S. macroeconomic aggregates, Vahid and Issler (2002) find a reduction of 20% for the MSFE of U.S. coincident series, and Athanasopoulos et al. (2011) find a reduction of 47% for the MSFE of Brazilian Inflation.

This paper confirms the usefulness of the common-feature VAR-based approach to forecast global measures of economic activity – global Industrial Production, U.S. GDP, and OECD's GDP. This paper also investigates possible forecast gains in economic activity by adding commodity prices to a baseline model, relying on the role commodity prices play in the derived-demand model. This is a standard use of commodity prices within the common-feature methodology. A different approach to this same issue benefits from the fact that one observes commodity prices on an almost continuous-time basis, but observe economic activity at a much lower frequency and with a substantial delay. Given the potential synchronicity of the cycles in economic activity and in commodity prices, one can employ a mixed-frequency (MIDAS, Andreou et al., 2013) framework to take advantage of the early signals (in chronological time) that commodity prices give for forecasting current economic activity. Notice that the advantage of the MIDAS approach relies on the fact that commodity prices are traded daily in global markets, while economic-activity data take time to be released by official authorities in their respective countries. This is the second original contribution of this paper.

The third original contribution of this paper is to compute optimal forecast solutions for oil prices, in the context of forecast combinations, where the number of forecasts, $i = 1, 2, \cdots, N$, and of time-series observations, $t = 1, 2, \cdots, T$, grow without bound; see Issler and Lima (2009). In this paper, the optimal forecast of oil-price growth rates at a given forecast horizon is the common feature of forecasts made by a diverse group of econometric models. The natural panel-data structure allows the identification and estimation of the oil-price optimal forecasts where the econometrician employs a mean-squared-error *risk function*. In this context, forecast combinations work well because of risk diversification: idiosyncratic forecast errors vanish since the law-of-large numbers works to eliminate the uncertainty associated with them as the number of forecasts being combined increases without bounds. Using these techniques, oil prices are forecasted conditioning on a large group of macroeconomic and financial predictors available from different databases: Goyal and Welch (2008), FRED-MD (McCracken and Ng, 2015), EPU (Economic Policy Uncertainty indexes of Baker, Bloom and Davis, 2015), GPR (Geopolitical Risk indexes of Caldara and Iacoviello, 2018).

The main focus is on the period from 1990 onwards for estimation, testing, and forecasting, because of its specificities. This period has been well known for an increase in the trade and financial openness of the global economy, something that has been labelled by some authors as *The Washington Consensus*, and by others simply as *Globalization.* Since 1990, there has been the emergence of important players in industrial production at the global level, such as China, India, and South Korea, as well as the emergence of resource-rich countries, that are commodity exporters, such as Brazil, Indonesia, Russia, and South Africa. Countries in these two groups are connected by a global supply chain.[2] Moreover, the previous decade (1980's) has witnessed the appearance of crude oil futures (West Texas Intermediate – WTI), as is known. Oil was first traded on the New York Mercantile Exchange (NYMEX) on March 30, 1983, but a consolidated oil market in size and scope of financial products is a more recent phenomenon dating from the 1990s.

First, the empirical results confirm the importance of a global demand factor for oil-price variations that are linked to the derived-demand model for oil. From 1990 onwards, Brent-price variations are synchronized with variations of global Industrial Production and of U.S. Industrial Production at the monthly and quarterly frequencies. At the quarterly frequency, Brent-price variations are also synchronized with variations of China's GDP. If one increases the sample period to include pre-1990 data, synchronization vanishes, which potentially highlights the importance of globalization to explain this phenomenon.

Second, it is shown that employing a VAR-based system approach consistent with the derived-demand model helps forecasting either global Industrial Production, U.S. GDP and OECD GDP. At the quarterly frequency, VAR-based forecasts that impose the restrictions consistent with the derived demand model for oil or for metal prices produce a smaller out-of-sample MSFE for U.S. GDP, compared to models that ignore the link between commodity prices and global economic activity. Results for OECD's GDP are also positive in general, although on occasion the opposite is observed. The

[2] For the record, focusing on a longer time span one does not find here the same evidence of synchronization found for the period 1990 onwards. This evidence indicates the synchronization of global output and crude-oil prices is a recent phenomenon linked to economic globalization and to the appearance of a proper trading market for oil during the 1980's. Of course, this is just a conjecture from our part. Further exploration of this topic is beyond the scope of this paper.

MIDAS approach relying on common cycles between the growth rate of oil prices and of global measures of economic activity improves predicting U.S. GDP growth rates, but not OECD's GDP growth rates.

Finally, the forecast combination approach for Brent-price variations performs best when compared with commonly used alternatives, like the weighted-average forecast using the reciprocal of the MSFE as weights, the best model chosen by BIC, and the random-walk model. In particular, it beats the random-walk model at different horizons – with a maximum out-of-sample gains of about 14%.

The paper is divided as follows: Section 2 presents a theoretical model that delivers common cycles among oil prices and industrial output. Section 3 summarizes the econometric techniques employed here. A complete treatment is presented on the Online Appendix to this paper. The empirical results are reported in detail in Section 4. Section 5 concludes.

2 Understanding the Short Run Fluctuations of Oil and Metal Prices

It is argued here there is an important role for demand shocks in explaining the short-run variation of oil and metal prices. These are important globally traded commodities, whose supply function is very price inelastic in the short run – zero or close to zero elasticity. The reason for that lies on the fact that the production function of these commodities is very capital intensive and it takes time to build capital. The Online Appendix to this paper has a more complete treatment, based on Issler et al. (2014).

Consider a representative cost-minimizing global industrial firm, which chooses the optimal quantity of inputs x_i, $i = 1, 2, \cdots, n$, all stacked in a vector $x = (x_1, x_2, \cdots, x_n)'$, when producing output y_0. The choice of output y_0 can be thought as an optimal decision coming from the firm's output market. The corresponding prices for inputs $i = 1, 2, \cdots, n$, stacked in a vector $w = (w_1, w_2, \cdots, w_n)'$, are considered given for the firm when choosing x. The firm's cost minimization problem in this context is:

$$\min_{x} \quad C(w, x) = w \cdot x \qquad s.t. \qquad f(x) \geq y_0. \tag{1}$$

If the supply of commodity i is given in the short run, an increase in y_0 will lead to an increase in w_i; see the Online Appendix. This result is completely intuitive: given concavity of the cost function, if the representative firm wants to increase industrial production in the short run, it will put an upward pressure on the price of these commodities, stemming from the fact that it should take more inputs to produce more, otherwise the firm is not a cost minimizer.

Next, there is a simple theoretical example illustrating this result, where production follows the well-known Cobb-Douglas production function with constant returns to scale:

$$y = Ax_1^{\alpha}x_2^{1-\alpha}, \qquad 0 < \alpha < 1. \tag{2}$$

The first input (input 1) will represent the input with a short-run fixed supply (oil, oil derivatives, and metal commodities). The Lagrangian for the cost-minimizing firm reads as:

$$\mathcal{L} = w_1x_1 + w_2x_2 - \lambda\left(y_0 - Ax_1^{\alpha}x_2^{1-\alpha}\right),$$

where λ is the Lagrange Multiplier. The first-order conditions are:

$$w_1 = \lambda\alpha A\left(\frac{x_2}{x_1}\right)^{1-\alpha}, \tag{3}$$

$$w_2 = \lambda(1-\alpha)A\left(\frac{x_2}{x_1}\right)^{-\alpha}, \tag{4}$$

$$y_0 = Ax_1^{\alpha}x_2^{1-\alpha}. \tag{5}$$

Using equations (3)-(5), it is straightforward to find the optimal demand for inputs 1 and 2 are respectively:

$$x_1^*(w_1,w_2,y_0) = \frac{y_0}{A}w_2^{1-\alpha}w_1^{\alpha-1}\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha},$$

$$x_2^*(w_1,w_2,y_0) = \frac{y_0}{A}w_1^{\alpha}w_2^{-\alpha}\left(\frac{1-\alpha}{\alpha}\right)^{\alpha}.$$

If one assumes a short-run fixed supply for input 1 (say, $\overline{x_1^s}$), then, its short-run market equilibrium condition is:

$$\overline{x_1^s} = \frac{y_0}{A}w_2^{1-\alpha}w_1^{\alpha-1}\left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha}. \tag{6}$$

Taking logs of (6), solving for $\ln(w_1)$ gives:

$$\ln(w_1) = \frac{1}{1-\alpha}\ln(y_0) - \frac{1}{1-\alpha}\ln\left(\overline{x_1^s}\right) - \frac{1}{1-\alpha}\ln(A) + \ln(w_2) + \ln\left(\frac{\alpha}{1-\alpha}\right),$$

where $\ln\left(\overline{x_1^s}\right)$ and $\ln(A)$ are constants, but, $\ln(w_2)$ depends on $\ln(y_0)$, through an equilibrium condition analogous to (6). As long as input 2 supply responds proportionally less than demand to a change in $\ln(y_0)$, one should have $\frac{\partial\ln(w_2)}{\partial\ln(y_0)} > 0$, obtaining a positive short-run elasticity and the synchronicity of cycles for $\ln(w_1)$ and $\ln(y_0)$:

$$\frac{\partial\ln(w_1)}{\partial\ln(y_0)} = \frac{1}{1-\alpha} + \frac{\partial\ln(w_1)}{\partial\ln(w_2)}\frac{\partial\ln(w_2)}{\partial\ln(y_0)} = \frac{1}{1-\alpha} + \frac{\partial\ln(w_2)}{\partial\ln(y_0)} > 0.$$

3 Econometric Techniques

The VAR- and the GMM-based Approaches

This section presents a summarized version of the econometric techniques employed in this paper. A more complete treatment is presented in the Online Appendix, which is based on Issler and Vahid (2001). Assume that y_t is a n-vector of $I(1)$ series that includes (log) oil prices and industrial production and additional relevant series, which can be well represented by a vector autoregression (VAR) model in levels.

One natural question that arises is how the theoretical model of the last section interacts with the econometric models employed here. As noted by King, Plosser, Stock and Watson (1991), in their well-known seminal use of VAR models, "Equation (5) [Wold Representation, consistent with a VAR] is a reduced-form relation and, except for purposes of forecasting, is of little inherent interest. What is of interest is the set of structural relations that leads to (5)..." From our perspective, a VAR model containing the economic series discussed in the last section is a reduced form, but theory allows one to use it to test for structural relations – the most relevant for this paper is the synchronicity of the commodity-price cycles and global economic activity, but one can also think of others involving long-run relationships (cointegration), and structural identification using orthogonalization procedures as well.

If the elements of y_t cointegrate, the VAR for y_t can be written as a vector error-correction model (VECM). Vahid and Engle (1993) show the VECM representation will be restricted if there exist white noise independent linear combinations $\tilde{\alpha}_i^{'}$, stacked in an $s \times n$ matrix $\tilde{\alpha}'$, of the series $\Delta\, y_t$, i.e., if the elements of $\Delta\, y_t$ share *common-cyclical* features. They consider an $\tilde{\alpha}$ matrix normalized as $\tilde{\alpha} = \begin{bmatrix} I_s \\ \tilde{\alpha}^*_{(n-s)\times s} \end{bmatrix}$ and complete the system by adding the last $n-s$ unconstrained VECM equations:

$$\begin{bmatrix} I_s & \tilde{\alpha}^{*\prime} \\ \underset{(n-s)\times s}{\mathbf{0}} & I_{n-s} \end{bmatrix} \Delta\, y_t = \begin{bmatrix} & \underset{s\times(np+r)}{\mathbf{0}} & & \\ \Gamma_1^{**} & \dots & \Gamma_{p-1}^{**} & \gamma^* \end{bmatrix} \begin{bmatrix} \Delta\, y_{t-1} \\ \vdots \\ \Delta\, y_{t-p+1} \\ \alpha' y_{t-1} \end{bmatrix} + v_t, \qquad (7)$$

where $v_t = \tilde{\alpha}'\epsilon_t$, $\alpha' y_{t-1}$ are the error-correction terms, and matrices $\Gamma_1^{**}, \Gamma_2^{**}, \dots, \Gamma_{p-1}^{**}$ and γ^* come from the last $n-s$ unconstrained VECM equations. Note the loading matrix of $\Delta\, y_t$ is always invertible, so one can recover a restricted reduced-form VECM for $\Delta\, y_t$ by pre-multiplying (7) by its inverse.

Estimation of the system in (7) can be performed by employing full-information maximum likelihood (FIML) under a known distribution for v_t, or by the generalized method of moments (GMM) if the properties of v_t are unknown. For the type of data considered in this paper, GMM is the preferred method. It exploits the fact that the errors v_t are orthogonal to the regressors in (7), $\alpha' y_{t-1}$, $\Delta\, y_{t-1}$, $\Delta\, y_{t-2}, \cdots, \Delta\, y_{t-p+1}$, collected in a vector Z_{t-1}:

$$\mathbf{0} = \mathbb{E}\left[v_t \otimes Z_{t-1}\right], \tag{8}$$

i.e., the orthogonality between all the elements in v_t and all the elements in Z_{t-1}. The test for common cycles is an over-identifying restriction test – the J test proposed in Hansen (1982) – which has an asymptotic χ^2 distribution with degrees of freedom equal to the number of over-identifying restrictions.

Optimal Forecast Combination

The common-feature approach in Issler and Lima (2009) also offers optimal forecasting techniques under a mean-squared-error *risk function*; see the Online Appendix for a broader treatment. They are appropriate for forecasting a weakly stationary and ergodic univariate process $\{y_t\}$ using a large number of forecasts that will be combined to yield an optimal forecast. These forecasts are the result of several econometric models that need to be produced prior to combining forecasts. Forecasts of y_t (change in log oil prices),[3] computed using conditioning sets lagged h periods, are denoted by $f_{i,t}^h$, $i = 1, 2, \ldots, N$. Therefore, $f_{i,t}^h$ are h-step-ahead forecasts and N is the number of models estimated to forecast y_t.

Issler and Lima assume forecasts $f_{i,t}^h$'s are modelled as approximations to the optimal forecast ($\mathbb{E}_{t-h}(y_t)$) as follows:

$$f_{i,t}^h = \mathbb{E}_{t-h}(y_t) + k_i^h + \varepsilon_{i,t}^h, \tag{9}$$

where k_i^h is the individual model time-invariant bias for h-step-ahead prediction and $\varepsilon_{i,t}^h$ is the individual model error term in approximating $\mathbb{E}_{t-h}(y_t)$, where $\mathbb{E}(\varepsilon_{i,t}^h) = 0$ for all i, t, and h. Here, the optimal forecast is a *common feature* of all individual forecasts and k_i^h and $\varepsilon_{i,t}^h$ arise because of forecast misspecification. The term k_i^h is identically distributed (but not independent), i.e., $k_i^h \sim i.d.\left(B^h, \sigma_{k^h}^2\right)$.

The *feasible* bias-corrected average forecast (BCAF) $\frac{1}{N}\sum_{i=1}^N f_{i,t}^h - \widehat{B^h}$, where $\widehat{B^h}$ is a consistent estimate of B^h, is an optimal forecasting device obeying:

$$\underset{(T,N\to\infty)_{seq}}{\operatorname{plim}} \left(\frac{1}{N}\sum_{i=1}^N f_{i,t}^h - \widehat{B^h}\right) = \mathbb{E}_{t-h}(y_t),$$

where $\underset{(T,N\to\infty)_{seq}}{\operatorname{plim}}$ is the probability limit using the sequential asymptotic framework of Phillips and Moon (1999). Issler and Lima also propose a test for zero bias, i.e., $H_0 : B^h = 0$, using the approach in Conley (1999). In case H_0 is not rejected in this test, there is no need to use a bias correction device in forecasting. Therefore, the optimal forecast will be the simple cross-sectional average of forecasts: $\frac{1}{N}\sum_{i=1}^N f_{i,t}^h$.

[3] Forecast-accuracy statistics can be based on log oil prices instead.

MIDAS Regressions

Following the notation of Andreou et al. (2013), assume the variable of interest is observed at some low frequency (e.g., real GDP quarterly growth rate), denoted by Y_t^Q, and the goal is to forecast this variable $h-$quarters ahead, that is Y_{t+h}^Q. The econometrician has daily observations of the financial predictors X. These are the growth rates of commodity prices. Denote by $X_{m-j,t}^D$ the jth day counting backward in quarter t, where m denotes the number of trading days per quarter – assumed to be constant for the sake of simplicity (e.g., $m = 66$). Hence, $X_{m,t}^D$, considering $j = 0$, corresponds to the observation of X on the last day of quarter t; see the Online Appendix for a broader treatment.

The ADL-MIDAS (p_Y^Q, q_X^D) model is given by:

$$Y_{t+h}^Q = \mu^h + \sum_{j=0}^{p_Y^Q-1} \rho_{j+1}^h Y_{t-j}^Q + \beta^h \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{m-1} \omega_{i+j*m}^{\theta^h} X_{m-i,t-j}^D + \varepsilon_{t+h}^h, \tag{10}$$

which entails p_Y^Q lags of Y_t^Q and q_X^D lags of $X_{m-i,t}^D$. The weighting scheme ω involves a low-dimensional vector of unknown parameters θ, used to avoid the parameter proliferation implied by the estimation of coefficients associated to high frequency lags. Following Andreou et al., the exponential Almon lag polynomial is adopted here.

To be consistent with the observed synchronicity of the cycles in GDP growth and on commodity prices represented by the $X_{m,t}^D$ variables, one must impose the restriction:

$$\rho_1^h = \rho_2^h = \cdots = \rho_{p_Y^Q}^h = 0, \tag{11}$$

since lagged GDP growth cannot enter in the white noise linear combination, thus making ε_{t+h}^h a white-noise process.

This setup in (10) can be extended to include macroeconomic data available at monthly frequency. Clements and Galvão (2008) introduced the MIDAS regression with *leads*, where the notion of *leads* pertains to the fact that one can use information between t and $t + 1$. Suppose one is 2 months into quarter $t + 1$ (i.e., at the end of February, May, August or November). This implies one has approximately 44 trading days (2 months) of daily data.

The ADL-MIDAS (p_Y^Q, q_X^D, J_X^M, J_X^D) regression model with *leads* (in both monthly and daily data) is described as follows (see Andreou et al., 2013, equation 5.1):

$$\begin{aligned} Y_{t+h}^Q \; &= \; \mu^h + \sum_{j=0}^{p_Y^Q-1} \rho_{j+1}^h Y_{t-j}^Q + \sum_{j=3-J_X^M}^{2} \gamma_j^h X_{3-j,t+1}^M \\ &\quad + \beta^h \left[\sum_{i=(3-J_X^D)*m/3}^{m-1} \omega_{i-m}^{\theta^h} X_{m-i,t+1}^D + \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{m-1} \omega_{i+j*m}^{\theta^h} X_{m-i,t-j}^D \right] + \varepsilon_{t+h}^h, \end{aligned} \tag{12}$$

where $X^D_{m-i,t+1}$ denotes the ith day counting backward in quarter $t+1$, and J^D_X represents daily *leads*, for the daily predictor, in terms of multiples of months (i.e., $J^D_X = 1$ and 2). In other words, in the case of $J^D_X = 2$, $X^D_{2m/3,t+1}$ corresponds to $2m/3 = 44$ leads (assuming $m = 66$), while $X^D_{1,t+1}$ corresponds to 1 lead for the daily predictor. The definitions of J^M_X and $X^M_{3-j,t+1}$, associated to monthly data, are quite similar to the ones used for daily data. Again, to have a MIDAS model consistent with the synchronicity of cycles between GDP and commodity prices one must impose (11).

4 Empirical analysis

The focus is on the period from 1990 onwards because of its specificities. This period has been well known for an increasing trade and financial openness of the global economy, something that has been labelled by some authors as *Globalization.* Since crude oil is a commodity that has a global character and has been traded in financial markets from the 1980s onwards, but *Globalization* is a 1990s phenomenon, the focus on this period makes sense. For the record, focusing on a longer time span one does not find here the same evidence found for the period 1990 onwards.[4]

4.1 Data and Empirical Implementation

On a monthly basis, the price of Brent Crude Oil and global industrial production series were both extracted from the International Financial Statistics (IFS) of the IMF. At this frequency they are available from 1990:M1 through 2019:M5. Nominal oil price data were deflated using the producer price index (PPI) for the U.S., which was extracted from the FRED database of the St. Louis FED.

On a quarterly basis, data for Brent Crude Oil, available from 1990:Q1 through 2019:Q1, were extracted from IFS database. Nominal price data were deflated using the PPI for the U.S. Metal-price data are also available from the IFS database for the period from 1990:Q1 through 2019:Q1. Global industrial production at the quarterly frequency was computed as the average monthly series. The U.S. GDP data is also obtained from the FRED database, kept by the St. Louis FED. The Chinese and Global GDP data were obtained from the IFS database, from 1990:Q1 through 2019:Q1, and the OECD GDP data were extracted from the OECD database for the same period.

In the forecasting exercise, using model combination, covariates (predictors) which are potentially correlated with the price of oil were employed. The start is the list of covariates in Hong and Yogo (2009, 2012). In addition to this data, out-of-sample forecasting exercises were conducted to select 20 monthly variables using the adaptive lasso approach of Zou (2006). These are potential covariates to forecast oil prices.

[4]Balcilar et al. (2017) find a similar result when they investigate the existence of common trends and cycles for oil prices and the S&P500 index using a long-span data set. Indeed, they only find a common cycle for the post-WW II period.

They came from a pool of 254 contemporaneous variables that are present in different databases: Goyal and Welch (2008), FRED-MD (McCracken and Ng, 2015), EPU (Economic Policy Uncertainty indexes of Baker, Bloom and Davis, 2015), GPR (Geopolitical Risk indexes of Caldara and Iacoviello, 2018) and the Thomson Reuters database. These 254 variables are lagged from one up to twelve months, forming a final large database containing 3,048 time series, from which those 20 series more frequently selected by the adaptive lasso were chosen, in an out-of-sample forecasting exercise of Brent oil price variation.

4.2 Cointegration and Common-Cycle Tests for Oil Prices

Monthly Frequency

Data for the (log) real prices of oil and for (log) U.S. Industrial Production (seasonally adjusted) are available from 1990:01 through 2019:05, whereas data for (log) Global Industrial Production (seasonally adjusted) are available from 1992:01 through 2019:05. All these series show signs of containing a unit root, which is confirmed for all of them using the Phillips and Perron (1989) test. Testing for stationarity rejected the null of stationarity using the Kwiatkowski et al. (1992) – KPSS test, which gives us confidence to model these series as $I(1)$.

First, this section investigates whether prices for oil cointegrate and/or share common cycles with U.S. and Global industrial production. Results are presented in Table 1. Regarding cointegration, there is no evidence of a long-run relationship between oil prices and U.S. and Global Industrial Production, respectively, for the last 29 years. On the other hand, results for common cycles are very different. Using a GMM approach, at the 5% significance level, there is evidence of strong-form *common-cyclical features* between these two measures of industrial production and oil prices: U.S. and Global Industrial Production, although results for the U.S. Industrial Production were borderline, near rejection by the *J-Statistic*. As a robustness check, this exercise also experimented with alternative lag-length choices in the VAR-based analysis, but the results were robust to alternative choices.[5] This section also tested for common cycles when the Kilian Index (Kilian, 2009) was used as a measure of global economic activity. Since this index is already in cyclical form (%, deviations from trend), the Beveridge and Nelson (1981) cycle is extracted from the (log) of Brent oil prices in testing for synchronicity. Results are borderline (at 5% significance) for our preferred specification with two lags, but using a three-lag structure confirms the existence of

[5] It is adopted here the standard procedure in the literature of inferring the lag lenght of the VAR using information criteria, test for cointegration, and then, conditional on cointegration, test for common cycles. As pointed out in Vahid and Issler (2002) and Athanasopoulos et al. (2011), there are some risks of implementing such strategy without checking whether the errors are indeed multivariate white noise in a context where short-run restrictions (common cycles) are likely to hold or where both short- and long-run restrictions (common cycles and cointegration) are likely to hold. Therefore, diagnostic tests are performed on VAR residuals to be on the safe side.

common cycles between Brent prices and the Kilian Index.

Table 1 - Common-Feature Tests (monthly 1991M01-2019M05)

$\Delta y_{1,t}$ (1,	$\Delta y_{2,t}$ α^*)	α^*	J-statistics
Oil	US Industrial Production (SF) 4L	-5.08**	0.039
		(2.27)	{0.066}
Oil	Global Industrial Production (SF) 4L	-7.07***	0.026
		(1.88)	{0.284}
Oil	Global Industrial Production (WF) 4L	-7.36***	0.027
		(1.86)	{0.344}
Oil	Kilian Index (SF) 2L	-0.366*	0.022
		(0.19)	{0.062}
Oil	Kilian Index (SF) 3L	-0.39**	0.026
		(0.19)	{0.115}

Robust Standard Errors (HAC) are in parentheses and p-values are in braces. 2L, 3L and 4L means a VAR with 2, 3 and 4 lags. SF and WF denote, respectively, strong- and weak-form common-cyclical features tests. Testing the Kilian Index used the Beveridge and Nelson (1981) cycle of (log) Brent price, not its first difference. The Kilian Index was divided by 100, since it measured in percentages. *,**, and *** represent significance at the 10%, 5% and 1% levels.

Figure 1 below plots the growth rates of oil prices – labelled $\Delta \ln\left(p_t^{OIL}\right)$ – and the growth rates of global industrial production – labelled $\Delta \ln\left(ip_t\right)$, both standardized (zero mean, unit variance). These series show clear signs of serial correlation. But, at the 5% significance level, the empirical results in Table 1 found that the following linear combination is white noise (unpredictable)[6]:

$$\Delta \ln\left(p_t^{OIL}\right) - \underset{(1.86)}{7.07} \times \Delta \ln\left(ip_t\right) \underset{(0.006)}{+0.013}, \tag{13}$$

confirming that $\Delta \ln\left(p_t^{OIL}\right)$ and $\Delta \ln\left(ip_t\right)$ are synchronized and that $\ln\left(p_t^{OIL}\right)$ and $\ln\left(ip_t\right)$ share a common cycle.

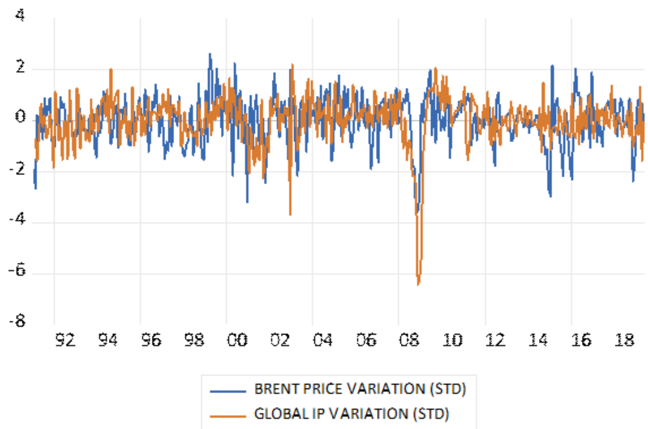

Figure 1: Monthly Growth Rates of Brent Prices and Global Industrial Production (standardized)

[6]Equation (13) represents the residual of the first equation in the system (7).

Quarterly Frequency

On a quarterly frequency, data for (log) oil prices are available from 1990:Q1 through 2019:Q1. Again, the nominal series is deflated by U.S. PPI. This section investigates the existence of pairwise *common features* between oil prices and aggregate measures of (log) output, such as Industrial Production (U.S. and Global), U.S. GDP, and China's GDP. All these series show signs of containing a unit root, which is confirmed for all of them using the Phillips-Perron test and the KPSS test.

There is no cointegration between real oil prices and these output measures. On short-run relationships, Table 2 presents results of pairwise *common-cycle* tests between oil prices and these output measures. Using the *J-statistic*, at the 5% level, there is overwhelming evidence of a strong-form common cycle between the instantaneous growth rate of oil prices and that of U.S. GDP, U.S. Industrial Production, and Global Industrial Production. Evidence for China's GDP is not very robust. Notice that different choices for the lag length of the VAR make no difference for the existence of common cycles. Focusing on the significance (strength) of the white-noise linear combination, as well at the p-value of the *J-statistic*, there is stronger evidence of a common cycle for oil prices and Global Industrial Production, and to a lesser extent for oil prices and U.S. Industrial Production, China's GDP, and U.S. GDP, on that order. Tests using the Kilian Index do not show conclusive results, since the coefficient on the index was not significant at usual levels, despite the fact that the over-identifying-restriction test did not reject the null of orthogonality between errors and instruments.

Table 2 - Common-Feature Tests Quarterly (1990Q01-2019Q01)

$\Delta y_{1,t}$	$\Delta y_{2,t}$	SCCF α^*	J-statistic
Oil	US Industrial Production (SF) 1L	-5.67*	0.013
		(3.11)	{0.213}
Oil	US Industrial Production (SF) 2L	-6.38**	0.023
		(2.82)	{0.436}
Oil	US GDP (SF) 3L	-11.28	0.042
		(6.92)	{0.422}
Oil	US GDP (SF) 4L	-5.40	0.061
		(5.99)	{0.621}
Oil	China GDP (SF) 4L	-2.51	0.041
		(2.89)	{0.707}
Oil	China GDP (WF) 4L	-8.80**	0.048
		(4.06)	{0.722}
Oil	Global Industrial Production (SF) 2L	-4.55**	0.027
		(1.91)	{0.397}
Oil	Global Industrial Production (SF) 3L	-4.45**	0.037
		(1.78)	{0.549}
Oil	Kilian Index (SF) 1L	-0.16	0.003
		(0.27)	{0.589}
Oil	Kilian Index (SF) 2L	-0.31	0.020
		(0.23)	{0.536}

Robust Standard Errors (HAC) are in parentheses and p-values are in braces.
1L, 2L, 3L and 4L means a VAR with 1, 2, 3 and 4 lags. SF and WF denote, respectively, strong- and weak-form common-cyclical features tests. Testing the Kilian Index used the Beveridge and Nelson (1981) cycle of (log) Brent price, not its first difference. The Kilian Index was divided by 100, since it measured in percentages. *,**, and *** represent significance at the 10%, 5% and 1% levels.

As an example of our findings, $\Delta \ln \left(p_t^{OIL}\right)$ and $\Delta \ln \left(ip_t\right)$ (global IP) show clear signs of serial correlation. However, at the usual 5% or 10% significance levels, the empirical results in Table 2 found that the following linear combination is white noise:

$$\Delta \ln \left(p_t^{OIL}\right) - \underset{(1.784)}{4.45} \times \Delta \ln \left(ip_t\right) + \underset{(0.0167)}{0.0146} . \tag{14}$$

4.3 Forecasting Global Economic Activity using Commodity Prices

A large amount of studies investigate the relationship between oil prices and economic activity. In general, they point out oil price shocks as recurrent sources of economic fluctuations, and often suggest that a positive innovation in economic activity is connected with an oil price increase; see Ratti and Vespignani (2016) and Bjørnland, Larsen and Maih (2018). In particular, Narayan et al. (2014) test whether oil prices predict economic growth, providing empirical evidence of greater predictability for developed countries. Herrera, Lagalo and Wada (2015) evaluate the presence of asymmetries in the relationship between economic activity and oil prices for a set of OECD countries, containing both oil exporters and oil importers. In turn, Hamilton (1983), Kilian (2009), and Kilian and Vigfusson (2017) investigate the extent to which oil price shocks explain U.S. recessions. Regarding oil shocks, macroeconomic fundamentals and monetary policy, Hamilton and Herrera (2004) challenge the view that monetary policy could be used to eliminate adverse consequences of oil shocks.

This section verifies the usefulness of the common-feature VAR-based approach for forecasting global measures of economic activity – U.S. GDP, global Industrial Production (IP), and OECD's GDP. Commodity prices are employed to help forecasting economic activity, relying again on the synchronicity between the cycles of these commodities and measures associated with global economic activity.

The proposed forecasting experiment entails the following steps:

1. Build a model only containing real economic activity, such as GDP (U.S., or OECD's) and global industrial production. This model can potentially exploit common trends and common cycles among different economic activity variables, with no role for commodity prices.

2. Build a model as in (1), augmented with commodity prices (oil or metal prices), which can additionally exploit the synchronicity of cycles generated by the derived-demand model.

3. Do a pseudo out-of-sample forecasting for these models separately, comparing their forecast accuracy, taking the model in (1) as the benchmark.

This section performs two types of experiments. In the first, all data are included in the same frequency, simply exploiting the parsimony that results in having a model with short-run restrictions, i.e., equation (7), but augmented with commodity prices.

The latter is compared with a model that could potentially have short-run restrictions but that does not include commodity prices. In the second experiment, it is considered the fact that commodity prices are sampled at a frequency higher than the one GDP data are sampled: daily versus quarterly. Therefore, exploiting synchronicity at the monthly frequency one can improve quarterly GDP forecasts by employing a mixed-frequency MIDAS approach as in Andreou et al. (2013).

The model including only quarterly real U.S. GDP growth rates $\left(\Delta \ln y_t^{US}\right)$ and global IP growth rates $(\Delta \ln ip_t)$ (model in (1)) exploits the existence of common cycles between these two series. The cofeature vector eliminating common serial correlation is $\widetilde{\alpha}' = (1, 0.426)$.

The full-sample *quasi-structural* model estimates for $\left(\Delta \ln y_t^{US}\right)$ and $(\Delta \ln ip_t)$ reads as:

$$\begin{bmatrix} 1 & \underset{(0.333)}{0.426} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln y_t^{US} \\ \Delta \ln ip_t \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \Delta \ln y_{t-1}^{US} \\ \Delta \ln ip_{t-1} \end{bmatrix} + 2nd\ Lag$$

with standard errors in parentheses, the stars representing non-zero coefficients which are omitted to save space. The *J-test Statistic* p-value is 0.3410, not rejecting the null of orthogonality between errors and instruments. Note, however, that the coefficient of $(\Delta \ln ip_t)$ in the common-cycle combination is not significant at the usual levels.

Issler et al. (2014) show evidence of common cycles between metal prices and global industrial production. The VECM model including metal-price growth rates (Aluminum, $\Delta \ln p_t^{Al}$, Copper, $\Delta \ln p_t^{Co}$, and Tin, $\Delta \ln p_t^{Tin}$) and their respective common-cycle restrictions with respect to global industrial production is presented below, omitting less-important parameter estimates. These three metal prices are the ones that share a common cycle with global industrial production, as documented in Issler et al. The *J-test Statistic* p-value for the system is 0.9998, not rejecting the null of orthogonality between errors and instruments. This model is a good example of the parsimony entailed by the existence of common-cyclical features.

$$\begin{bmatrix} 1 & 0 & 0 & 0 & \underset{(0.369)}{-5.640} \\ 0 & 1 & 0 & 0 & \underset{(0.789)}{-5.587} \\ 0 & 0 & 1 & 0 & \underset{(0.597)}{-4.129} \\ 0 & 0 & 0 & 1 & \underset{(0.021)}{-0.232} \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln p_t^{Al} \\ \Delta \ln p_t^{Co} \\ \Delta \ln p_t^{Tin} \\ \Delta \ln y_t^{US} \\ \Delta \ln ip_t \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * \end{bmatrix} \begin{bmatrix} \Delta \ln p_{t-1}^{Al} \\ \Delta \ln p_{t-1}^{Co} \\ \Delta \ln p_{t-1}^{Tin} \\ \Delta \ln y_{t-1} \\ \Delta \ln ip_{t-1} \end{bmatrix}$$

$$+ \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \underset{(0.003)}{-0.031} & \underset{(0.002)}{0.017} \end{bmatrix} \underbrace{EC_{t-1}}_{2\times 1},$$

where EC_{t-1} is the 2×1 error-correction term in the VECM.

Regarding the MIDAS approach, this section forecasts the real GDP growth rate (Y^Q_{t+h}) for the U.S. or the OECD, $h-$quarters ahead. It is used as a monthly variable $X^M_{month=i} = \Delta \ln(\text{global industrial production index}_{month=i})$, and as daily variables $X^D_{day=j} = \frac{1}{15}\sum_{k=1}^{15} \Delta \ln(\text{commodity price}_{day=j-k+1})$, where the set of commodities is either $[Oil]'$ or $[Aluminium; Copper; Tin]'$. The source of daily data is Thomson Reuters[7]. Regarding model (10), the predictors $[X^D; X^M]'$ are considered lagged two months in respect to the end of the quarter of the observed GDP (Y^Q_t). In model (12), the predictors $[X^D; X^M]'$ are considered one and two months forwarded, respectively, in respect to the end of the quarter of the observed GDP (Y^Q_t). In other words, for model (12), information between periods t and $t+1$ is taken into account when considering *leads* of one and two months, respectively, for the commodity prices and the global industrial production index. Therefore, one conditioned on a different information set of the other models entertained in the experiment.

Next, the results of the forecast exercise devised above for metal prices are presented. Forecast accuracy is measured using the root-mean-squared error (RMSE) in forecasting the logarithm of U.S. GDP $\left(\ln\left(y^{US}_t\right)\right)$ at different quarterly horizons, up to one year ahead. Table 3 presents the results of the forecasting experiment comparing the two models in (1) and (2) above, as well as with the two versions of the MIDAS models – ADL-MIDAS and ADL-MIDAS *with leads*. Overall, one can see the benefits of using metal prices in the system, especially at the short horizons for the VAR-based approach, and at all horizons for the ADL-MIDAS with leads, albeit not always significantly better.

Table 3: Forecasting U.S. GDP

RMSE ($\times$100) with and without Metal Prices

Model/Horizon	$h=1$	$h=2$	$h=3$	$h=4$
(1) System Without Metal Prices	0.438	0.639	0.809	1.055
(2) System With Metal Prices	0.429	0.594	0.810	1.118
(3) ADL-MIDAS	0.410	0.586	0.750	0.908*
(4) ADL-MIDAS with leads	0.377**	0.573	0.698*	0.867***

Notes: First estimation sample: 1991Q1-2009Q4 (evaluation sample: 2010Q1-2019Q1). Rolling window estimation. RMSE with $(38-h)$ out-of-sample observations. Model (1) is the benchmark. *,**,*** Denote significance at the 10%, 5% and 1% levels, respectively, using the Giacomini and White (2006) predictive ability test, indicating that model (i) is better than model (1).

The VECM model including oil-price growth rates $(\Delta \ln p^{OIL}_t)$ and its respective common-cycle restrictions with respect to global industrial production is presented

[7]Nominal price data are deflated using the producer price index (PPI) for the U.S., extracted from the FRED database of the St. Louis FED.

below. Note that the cofeature white noise linear combination $(1, 0, -13.91)'$ is rather different than the one presented in Table 2. The *J-test Statistic* p-value is 0.2949, showing no signs of model misspecification for this system.

$$\begin{bmatrix} 1 & 0 & \underset{(5.92)}{-13.91} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln p_t^{OIL} \\ \Delta \ln y_t^{US} \\ \Delta \ln ip_t \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix} \begin{bmatrix} \Delta \ln p_{t-1}^{OIL} \\ \Delta \ln y_{t-1}^{US} \\ \Delta \ln ip_{t-1} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix} \begin{bmatrix} \Delta \ln p_{t-2}^{OIL} \\ \Delta \ln y_{t-2}^{US} \\ \Delta \ln ip_{t-2} \end{bmatrix}.$$

Next, the results of the forecast exercise devised above for oil prices are shown. Table 4 presents the results of the forecasting experiment comparing the two models in (1) and (2) above, as well as with the two versions of the MIDAS models. Including oil prices in the system improves short-run forecasting economic activity, since it reduces out-of-sample RMSE, up to two quarters ahead for the VAR-based forecasts. However, this reduction is not significant using the Giacomini and White (2006) test. Results for the two MIDAS models show a further improvement in forecast accuracy, which are significant for all but one horizon.

These forecasting results in Table 4 are in line with the previous literature, since there is a long tradition of using oil prices to forecast U.S. real GDP. For example, Ravazzolo and Rothman (2012) investigate Granger-causal relationships between crude oil prices and U.S. GDP growth and compare a benchmark model without oil against alternatives including oil. The results indicate a strong rejection of the null hypothesis of no out-of-sample predictability from oil prices to GDP from the mid-80s through the Great Recession. In another example, Kilian and Vigfusson (2013) point out the role of nonlinearities and asymmetries in out-of-sample accuracy of real GDP growth forecasts using oil prices with focus on iterated forecasts (rather than direct forecasts).

Table 4: Forecasting U.S. GDP

RMSE (×100) with and without Oil Prices

Model/Horizon	$h = 1$	$h = 2$	$h = 3$	$h = 4$
(1) System Without Oil Prices	0.438	0.639	0.809	1.055
(2) System With Oil Prices	0.437	0.625	0.892	1.306
(3) ADL-MIDAS	0.381**	0.567	0.654**	0.798***
(4) ADL-MIDAS with leads	0.375**	0.563	0.697**	0.820***

Notes: See Table 3.

The focus now is on forecasting OECD's GDP. It is hard to find reliable global measures of economic activity going back to the early 1990s on a quarterly frequency. The search led us to OECD's GDP, which covers a wide array of countries in the developed world, but does not include emerging economies, which may be a problem, since most global industrial production is currently made in emerging countries. The step (1) model described above is:

$$\begin{bmatrix} 1 & \underset{(0.078)}{-0.510} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln y_t^{OECD} \\ \Delta \ln ip_t \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ * & * \end{bmatrix} \begin{bmatrix} \Delta \ln y_{t-1}^{OECD} \\ \Delta \ln ip_{t-1} \end{bmatrix},$$

where $\Delta \ln y_t^{OECD}$ is the instantaneous growth rate of OECD's GDP. The *J-test Statistic* p-value is 0.4278, not rejecting the null of orthogonality between errors and instruments. The model in step (2) including Brent oil prices is:

$$\begin{bmatrix} 1 & \underset{(1.754)}{-3.054} & 0 \\ 0 & 1 & \underset{(0.213)}{-1.784} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln p_t^{OIL} \\ \Delta \ln ip_t \\ \Delta \ln y_t^{OECD} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ * & * & * \end{bmatrix} \begin{bmatrix} \Delta \ln p_{t-1}^{OIL} \\ \Delta \ln ip_{t-1} \\ \Delta \ln y_{t-1}^{OECD} \end{bmatrix},$$

which again illustrates the potential efficiency gains that could be present when a model in *pseudo-structural* form is entertained. The *J-test Statistic* p-value is 0.2646, showing no signs of misspecification for the proposed structure.

Table 5 presents the results of the forecasting experiment comparing the two VAR-based models, as well as the two MIDAS models. Including oil prices in the system improves forecasting OECD GDP up to three quarters ahead. This reduction is not significant vis-a-vis the basic model at usual levels using the Giacomini and White (2006) test. The results for the two MIDAS models are disappointing for all horizons.

Table 5: Forecasting OECD's GDP

RMSE (×100) with and without Oil Prices

Model/Horizon	$h = 1$	$h = 2$	$h = 3$	$h = 4$
(1) System Without Oil Prices	0.229	0.390	0.526	0.641
(2) System With Oil Prices	0.214	0.367	0.515	0.649
(3) ADL-MIDAS	0.283	0.443	0.532	0.680
(4) ADL-MIDAS with leads	0.238	0.496	0.666	0.795

Notes: See Table 3.

The model containing metal prices and OECD's GDP reads as:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & \underset{(0.489)}{-5.537} \\ 0 & 1 & 0 & 0 & \underset{(0.819)}{-5.385} \\ 0 & 0 & 1 & 0 & \underset{(0.695)}{-4.488} \\ 0 & 0 & 0 & 1 & \underset{(0.051)}{-0.470} \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta \ln p_t^{Al} \\ \Delta \ln p_t^{Co} \\ \Delta \ln p_t^{Tin} \\ \Delta \ln y_t^{OECD} \\ \Delta \ln ip_t \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ * & * & * & * & * \end{bmatrix} \begin{bmatrix} \Delta \ln p_{t-1}^{Al} \\ \Delta \ln p_{t-1}^{Co} \\ \Delta \ln p_{t-1}^{Tin} \\ \Delta \ln y_{t-1}^{OECD} \\ \Delta \ln ip_{t-1} \end{bmatrix},$$

where the *J-test Statistic* p-value is 0.7284, not rejecting the null of orthogonality between errors and instruments. Table 6 presents the results of the forecasting experiment comparing these two models and the MIDAS models. Here, using the VAR-based approach, including metal prices in the system improves forecasting economic activity for horizons one and two. The MIDAS approach is also disappointing for forecasting OECD's GDP.

Table 6: Forecasting OECD's GDP				
RMSE (×100) with and without Metal Prices				
Model/Horizon	$h = 1$	$h = 2$	$h = 3$	$h = 4$
(1) System Without Metal Prices	0.229	0.390	0.526	0.641
(2) System With Metal Prices	0.226	0.385	0.540	0.671
(3) ADL-MIDAS	0.272	0.468	0.645	0.738
(4) ADL-MIDAS with leads	0.244	0.504	0.697	0.846
Notes: See Table 3.				

4.4 Forecasting Oil Prices using Forecast Combinations

In this section, the target variable in forecasting is the variation of the real price of Brent Crude Oil. So, results from this section are completely independent from the results presented in the last section. As explained above, a set of 3,048 covariate time series are put together coming from a list of contemporaneous 254 variables present in different databases: Goyal and Welch (2008), FRED-MD (McCracken and Ng, 2015), EPU (Economic Policy Uncertainty indexes of Baker, Bloom and Davis, 2015), GPR (Geopolitical Risk indexes of Caldara and Iacoviello, 2018) and also the Thomson Reuters database. To select useful covariates from that list, out-of-sample forecasting exercises for Brent oil price variation came up with 20 monthly variables, selected using the adaptive lasso approach of Zou (2006), shown below in Table 7:

Table 7 - Forecasting Covariates

Description	Source
United States Federal Funds Effective Rate.	Federal Reserve
United States Treasury Constant Maturity 10 Years.	Federal Reserve
3-Month Treasury C Minus FEDFUNDS	FRED-MD
Industrial Production - Total Iindex	Federal Reserve
Industrial Production - Manufacturing	Federal Reserve
S&P500 composite index	Reuters
CBOE SPX Volatility VIX	CBOE
Real Narrow Effective Exchange Rate for United States.	FRED
Primary metals leading index.	U.S.G.S.
Primary metals coincident index.	U.S.G.S.
CRB BLS Spot Index	Reuters
CRB BLS Spot Index Metals	Reuters
Geopolitical Risk Index in Ukraine	Caldara and Iacoviello
PPI: Metals and metal products	FRED-MD
IP: Business Equipment	FRED-MD
Production of Total Industry in Hungary	FRED
Industrial Production: Dur. Manuf.: Aerosp. and Misc. Transp. Equip	FRED
Industrial Production: Mining: Drilling oil and gas wells	FRED
Industrial Production: Durable Goods: Raw steel	FRED
Industrial Production: Durable Goods: Iron and steel products	FRED

In order to fit well the cross-sectional asymptotic requirement (large N) in the model-combination approach, one needs to have a large set of diversified forecasts to eliminate the combination of idiosyncratic errors. A few classes of different econometric models are chosen: ARMA and VAR models, all using distinct covariates (predictors), and distinct functional forms (levels, logs), and distinct stationarity assumptions (stationarity vs. difference-stationarity) for the target variable and predictors. Overall, this exercise uses $N = 32,412$ models.

The total sample, at the monthly frequency, comprises 341 time-series observations from 1991M01 until 2019M05. The sample is splitted in three distinct parts: data from 1 to T_1 (1991M01 to 2001M01) are used to estimate the coefficients of each model, containing 121 observations; data from $T_1 + 1$ to T_2 (2001M02 to 2011M01) are used to compute the bias, containing 120 observations; and data from $T_2 + 1$ to T (2011M02 to 2019M05) are used to implement pseudo out-of-sample forecasting, containing 100 observations.

To assess forecast accuracy, an algorithm appropriate for the bias-corrected average forecast (BCAF) is constructed. For alternative forecast combinations or forecasting schemes, slight modifications are required; see the Online Appendix for details. Out-of-sample R^2 statistics (percentage) comparing different forecast strategies with the random-walk model with or without drift are computed. The R^2-statistic for one-step ahead forecasts of the real price of oil z_t is:

$$R^2 = 100 \times \left[1 - \frac{\sum\limits_{t=T_2+1}^{T} \left(z_t - \widehat{z}_{t|t-1} \right)^2}{\sum\limits_{t=T_2+1}^{T} \left(z_t - \widehat{z}_{t|t-1}^{BMK} \right)^2} \right],$$

where $\widehat{z}_{t|t-1}$ is the one-step-ahead forecast of any given forecast strategy and $\widehat{z}_{t|t-1}^{BMK}$ is the one-step-ahead forecast of the benchmark – the random-walk model with or without a drift. Positive (negative) values for the R^2 statistic mean that the forecast $\widehat{z}_{t|t-1}$ beats (is beaten by) $\widehat{z}_{t|t-1}^{BMK}$.

Table 8, below, presents the result of the zero-bias test of Issler and Lima (2009), using the similarity of RMSFE across models to measure closeness in the cross section. The results show that, up to $h = 6$ months, all estimated biases are statistically nil. In this case, the optimal forecast in Issler and Lima's approach is simply the cross-sectional average forecast across models ($32,412$ of them).

Table 8: Mean bias significance test (1991M01 - 2019M05).

	Bias	t-stat	p-value
1 step-ahead	-0.84	-0.53	0.30
2 step-ahead	-1.62	-0.53	0.30
3 step-ahead	-2.28	-0.54	0.29
4 step-ahead	-2.95	-0.55	0.29
5 step-ahead	-3.55	-0.55	0.29
6 step-ahead	-4.07	-0.55	0.29

This table presents the results of the zero-bias test in Issler and Lima (2009).

Next, the out-of-sample forecast accuracy of different forecast strategies is examined: the random walk model (with or without drift), which is the benchmark, the "Best Model" chosen by the Bayesian Information Criterion (BIC) *ex-post*, the average of the best five (or ten) models, the average forecast across all models, and the Bias-Corrected-Average Forecast (BCAF) proposed by Issler and Lima.

Given the result of the zero-bias test above, the optimal forecast should be the simple (cross-sectional) average forecast. Indeed, looking at the results in Table 9 for out-of-sample forecast accuracy, this forecast strategy outperforms all other strategies from the one-month ahead to the six-month ahead horizon. If one takes the random walk model with drift as the benchmark, the out-of-sample R^2 statistics (percentage) vary from about 7% to about 14%, depending on the horizon.

Table 9: Out-of-Sample Forecast root-mean-squared-error and Out-of-Sample Forecast R^2 (1991M01 - 2019M05)

Horizon	BCAF	Weighted average (MSE)	Average forecast	Best model (BIC)	Average 5 best models	Average 10 best models
1 step-ahead	2.75	2.87	**2.50**	2.70	2.68	2.67
2 step-ahead	4.52	4.77	**4.11**	4.43	4.40	4.37
3 step-ahead	5.61	6.00	**5.14**	5.63	5.60	5.55
4 step-ahead	6.38	6.86	**5.84**	6.51	6.48	6.39
5 step-ahead	6.93	7.45	**6.29**	7.17	7.14	7.02
6 step-ahead	7.55	8.03	**6.83**	7.93	7.90	7.74
R^2 (RW w/drift 1-step ahead)	-8.30%	-17.87%	**10.06%**	-4.33%	-3.34%	-2.67%
R^2 (RW w/drift 2-step ahead)	-11.92%	-24.40%	**7.58%**	-7.17%	-5.99%	-4.68%
R^2 (RW w/drift 3-step ahead)	-10.68%	-26.39%	**7.04%**	-11.49%	-10.33%	-8.25%
R^2 (RW w/drift 4-step ahead)	-11.19%	-28.74%	**6.83%**	-15.75%	-14.59%	-11.60%
R^2 (RW w/drift 5-step ahead)	-7.31%	-23.88%	**11.72%**	-14.89%	-13.82%	-9.89%
R^2 (RW w/drift 6-step ahead)	-5.42%	-19.10%	**13.81%**	-16.27%	-15.34%	-10.72%
R^2 (RW w/o drift 1-step ahead)	-10.05%	-19.77%	**8.61%**	-6.02%	-5.01%	-4.33%
R^2 (RW w/o drift 2-step ahead)	-15.26%	-28.11%	**4.83%**	-10.37%	-9.15%	-7.81%
R^2 (RW w/o drift 3-step ahead)	-16.03%	-32.50%	**2.54%**	-16.88%	-15.66%	-13.48%
R^2 (RW w/o drift 4-step ahead)	-19.27%	-38.09%	**0.06%**	-24.15%	-22.91%	-19.71%
R^2 (RW w/o drift 5-step ahead)	-17.52%	-35.67%	**3.31%**	-25.83%	-24.66%	-20.35%
R^2 (RW w/o drift 6-step ahead)	-17.82%	-33.10%	**3.67%**	-29.94%	-28.90%	-23.74%

Bold values represents the best forecast. T = 341 (2019M05), T_1 = 120 (2000M12), T_2 = 240 (2010M12). BCAF: bias-corrected average forecast; MSE: mean squared error; the best model, the 5 best models, and 10 best models were chosen using the in-sample BIC criterium.

5 Conclusion and Further Research

This paper studies the interaction between commodity-price changes and global economic activity in a setup where a representative cost-minimizing firm chooses optimal inputs (commodity quantum) as derived demands. The focus is on important globally traded commodities, whose supply function is very price inelastic in the short run. Key examples studied in this paper are Oil (Oil derivatives) and also major metal commodities, such as Aluminum, Copper, Nickel.

The empirical evidence here fully supports the theoretical results of the derived-demand model. Indeed, this paper shows overwhelming evidence that cycles in oil prices are synchronized to those of global industrial production. This evidence is stronger regarding the global economy but holds as well for the U.S. economy. To the best of our knowledge, this is the first paper to investigate and find common cycles accounting for theory and empirics. This is the first original contribution of this paper.

The second original contribution is to investigate the usefulness of the common-feature VAR-based approach for forecasting global measures of economic activity – U.S. GDP and OECD's GDP. Here, the interest is reversed. Commodity prices are employed to help forecasting economic activity, relying again on the role commodities play in the derived-demand model. This is important, since one observes commodity prices on an almost continuous-time basis, but economic activity at a much lower frequency. So, commodity prices could give early signals about future global economic activity. For the most part, the empirical evidence in forecasting global measures of economic activity emphasizes the importance of conditioning on either oil prices or metal-commodity prices, relying on parameter-reduction techniques that favor optimal forecasts. Forecasting gains are substantial for U.S. GDP, but results for OECD's GDP are disappointing.

Finally, the third objective of the paper is to *forecast* oil prices over short horizons (monthly data) based on techniques that view the optimal forecast as a *common feature* (latent variable), which can be identified by using a cross-sectional average of a diverse group of forecasts. The model-combination approach considers 254 macroeconomic and financial covariates available from different databases: Goyal and Welch (2008), FRED-MD (McCracken and Ng, 2015), EPU (Economic Policy Uncertainty indexes of Baker, Bloom and Davis, 2015), GPR (Geopolitical Risk indexes of Caldara and Iacoviello, 2018). From this group, 20 covariates are selected and used on 32,412 different models – all predicting oil prices. When these models are combined using optimal techniques, they outperform the random-walk model with an out-of-sample R^2 statistic that can reach up to 13.81%. This is a major improvement vis-a-vis the previous literature.

References

[1] Aastveit, K.A., H.C. Bjørnland, and L.A. Thorsud, 2015, What drives Oil Prices? Emerging Versus Developing Economies? *Journal of Applied Econometrics* 30(1), 1013 –1028.

[2] Alquist, Ron, Kilian, Lutz, Vigfusson, Robert J., 2012. Forecasting the Price of Oil. Chapter Chapter 8 in *The Handbook of Economic Forecasting*, 2013, vol. 2, pp 427-507.

[3] Andreou, E., Ghysels, E., Kourtellos, A., 2013. Should Macroeconomic Forecasters Use Daily Financial Data and How? *Journal of Business and Economic Statistics* 31(2), 240-251.

[4] Athanasopoulos, G., Guillén, O.T.C., Issler, J.V., Vahid, F., 2011. Model Selection, Estimation and Forecasting in VAR Models with Short-run and Long-run Restrictions. *Journal of Econometrics* 164 (1), 116-129.

[5] Bai, J., Ghysels, E., Wright, J., 2013. State Space Models and MIDAS Regressions. *Econometric Reviews* 32(7), 779-813.

[6] Baker, S.R., Bloom, N., Davis, S.J., 2015. Measuring Economic Policy Uncertainty. NBER Working Paper 21633, National Bureau of Economic Research.

[7] Balcilar, M., Gupta, R., and Wohar, M.E, 2017. Common cycles and common trends in the stock and oil markets: Evidence from more than 150 years of data. *Energy Economics*, 61, pp. 72–86.

[8] Batchelor, R. (2007). Bias in macroeconomic forecasts. *International Journal of Forecasting* 23(2), 189 –203.

[9] Baumeister, C., Hamilton, J.D., 2019. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review* 109(5), 1873-1910.

[10] Baumeister, C., Kilian, L., Lee, T.K., 2014. Are there Gains from Pooling Real-Time Oil Price Forecasts? *Energy Economics* 46, S33-S43.

[11] Baumeister, C., Kilian, L., 2015. Forecasting the Real Price of Oil in a Changing World: A Forecast Combination Approach, *Journal of Business and Economic Statistics* 33(3), 338-351.

[12] Baumeister, C., Kilian, L., 2016. Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us, *Journal of Economic Perspectives* 30(1), 139-160.

[13] Benes, J., Chauvet, M., Kamenik, O., Kumhof, M., Laxton, D., Mursula, S., Selody, J., 2015. The Future of Oil: Geology versus Technology. *International Journal of Forecasting* 31(1), 207-221.

[14] Beveridge, S. and C. Nelson (1981). A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the 'business cycle'. *Journal of Monetary Economics* 7(2), 151–174.

[15] Bjørnland, H.C., Larsen, V.H., Maih, J., 2018. Oil and Macroeconomic (In)stability. *American Economic Journal: Macroeconomics* 10(4), 128-151.

[16] Caldara, D., Iacoviello, M., 2018. Measuring Geopolitical Risk. FRB International Finance Discussion Paper n. 1222, Board of Governors of the Federal Reserve Board.

[17] Clark, Todd E., West, Kenneth D., 2007. Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics* 138(5), 291–311.

[18] Clements, M.P., Galvão, A.B., 2008. Macroeconomic Forecasting With Mixed Frequency Data: Forecasting US Output Growth. *Journal of Business and Economic Statistics* 26, 546-554.

[19] Conley, T.G., 1999, GMM estimation with cross sectional dependence, *Journal of Econometrics*, 92(1), pp. 1-45.

[20] Deaton, A., Laroque, G., 1996. Competitive storage and commodity price dynamics. *Journal of Political Economy* 104 (5), 896–923.

[21] Dvir, Eyal and Kenneth Rogoff. 2014. "Demand Effects and Speculation in Oil Markets: Theory and Evidence." *Journal of International Money and Finance* 42: 113–128.

[22] Engle, R.F., Granger, C.W.J., 1987. Co-integration and error correction: representation, estimation, and testing. *Econometrica* 55(2), 251–276.

[23] Engle, R.F., Issler, J.V., 1995. Estimating common sectoral cycles. *Journal of Monetary Economics* 35 (1), 83–113.

[24] Engle, R.F., Kozicki, S., 1993. Testing for common features. *Journal of Business and Economic Statistics* 11 (4), 369–380.

[25] Gaglianone, W.P., and Issler, J.V., 2019, Microfounded Forecasting, Working Paper, FGV EPGE.

[26] Gargano, A., Timmermann, A., 2014. Predictive Dynamics in Commodity Prices. *International Journal of Forecasting*, vol. 30(3), pp. 825-843.

[27] Ghysels, E., Wright, J., 2009. Forecasting Professional Forecasters. *Journal of Business and Economic Statistics* 27, 504-516.

[28] Giacomini, R., White, H., 2006. Tests of conditional predictive ability. *Econometrica* 74, 1545-1578.

[29] Goyal, A., Welch, I., 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21(4), 1455-1508.

[30] Hamilton, J.D., 1983. Oil and the Macroeconomy since World War II, *Journal of Political Economy*, 91, 228-248.

[31] Hamilton, J.D., 2008. Daily Monetary Policy Shocks and New Home Sales. *Journal of Monetary Economics* 55, 1171-1190.

[32] Hamilton, J.D., Herrera, A., 2004. Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy. *Journal of Money, Credit, and Banking* 36(2), 265-286.

[33] Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50 (4), 1029–1054.

[34] Hecq, A., Palm, F., Urbain, J., 2006. Common cyclical features analysis in var models with cointegration. *Journal of Econometrics* 132 (1), 117–141.

[35] Herrera, A., Lagalo, L.G., Wada, T., 2015. Asymmetries in the Response of Economic Activity to Oil Price Increases and Decreases? *Journal of International Money and Finance* 50, 108-133.

[36] Herrera, A.M., Rangaraju, S.K., 2020. The effect of oil supply shocks on US economic activity: What have we learned? *Journal of Applied Econometrics* 35(2), 141-159.

[37] Hendry, D.F., Clements, M.P., 2004. Pooling of forecasts. *Economic Journal* 7 (1), 1–31.

[38] Hong, H., Yogo, M., 2009. Digging into Commodities. Unpublished Working Paper. Princeton University and Wharton of University of Pennsylvania.

[39] Hong, H., Yogo, M., 2012. What does futures market interest tell us about the macroeconomy and asset prices? *Journal of Financial Economics.* 105 (3), 473–490.

[40] IMF, 2012, April. World Economic Outlook: Growth Resuming, Dangers Remain. Chapter 1. *International Monetary Fund*, pp. 1–47.

[41] Issler, J.V., Lima, L.R., 2009. A panel data approach to economic forecasting: the bias-corrected average forecast. *Journal of Econometrics* 152 (2), 153–164.

[42] Issler, J.V., Rodrigues, C.F., and Burjack, R. (2014). Using Common Features to Understand the Behavior of Metal-Commodity Prices and Forecast them at Different Horizons. *Journal of International Money and Finance* 42, 310 –335.

[43] Issler, J.V., Vahid, F., 2001. Common cycles and the importance of transitory shocks to macroeconomic aggregates. *Journal of Monetary Economics* 47 (3), 449–475.

[44] Issler, J.V., Vahid, F., 2006. The missing link: using the nber recession indicator to construct coincident and leading indices of economic activity. *Journal of Econometrics* 132 (1), 281–303.

[45] Jerrett, D., Cuddington, J., 2008. Broadening the statistical search for metal price super cycles to steel and related metals. *Resources Policy* 33 (4), 188–195.

[46] Johansen, S., 1991. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica* 59 (6), 1551–1580.

[47] Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053-69.

[48] Kilian, L., 2019. Measuring global real economic activity: Do recent critiques hold up to scrutiny? *Economics Letters* 178, 106-110.

[49] Kilian, Lutz and Thomas K. Lee (2014). Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance* 42.

[50] Kilian, L., Murphy, D.P., 2012. Why agnostic sign restrictions are not enough: understanding the dynamics of oil market VAR models. *Journal of the European Economic Association* 10(5), 1166-1188.

[51] Kilian, L., Vigfusson, R.J., 2013. Do Oil Prices Help Forecast U.S. Real GDP? The Role of Nonlinearities and Asymmetries. J*ournal of Business and Economic Statistics* 31(1), 78-93.

[52] Kilian, L., Vigfusson, R.J., 2017. The Role of Oil Price Shocks in Causing U.S. Recessions. *Journal of Money, Credit, and Banking* 49(8), 1747-1776.

[53] Kilian, L., Zhou, X., 2018. Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance* 88 (2018) 54.78.

[54] King, R.G., Plosser, C.I., Stock, J.H., Watson, M.W., 1991. Stochastic Trends and Economic Fluctuations. *American Economic Review* 81, 819-840.

[55] Kydland, Finn E. and Edward C. Prescott, 1982, Time to Build and Aggregate Fluctuations, *Econometrica*, Vol. 50, No. 6, pp. 1345-1370.

[56] Kuzin, V., Marcellino, M., Schumacher, C., 2011. MIDAS Versus Mixed- Frequency VAR: Nowcasting GDP in the Euro Area. *International Journal of Forecasting* 21, 529-542.

[57] Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., Shin, Y. (1992). "Testing the null hypothesis of stationarity against the alternative of a unit root". *Journal of Econometrics*. 54 (1–3): 159–178.

[58] Laster, D., P. Bennett, and I. Geoum (1999). Rational bias in macroeconomic forecasts. *The Quarterly Journal of Economics* 114(1), 293–318.

[59] McCracken, M.W., Ng, S., 2015. FRED-MD: A Monthly Database for Macroeconomic Research. Working Paper 2015-012B, Federal Reserve Bank of St. Louis.

[60] Miller, J.I., Ni, S., 2011. Long-Term Oil Price Forecasts: A New Perspective on Oil and the Macroeconomy. *Macroeconomic Dynamics* 15(S3), 396-415.

[61] Miller, J.I., Ratti, R., 2009. Crude Oil and Stock Markets: Stability, Instability, and Bubbles. *Energy Economics* 31(4), 559-568.

[62] Narayan, P.K., Sharma, S., Poon, W.C., Westerlund, J., 2014. Do oil prices predict economic growth? New global evidence. *Energy Economics* 41, 137-146.

[63] Newey, W. and K. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 703–708.

[64] Palm, F. C. and A. Zellner (1992). To combine or not to combine? issues of combining forecasts. *Journal of Forecasting* 11(8), 687–701.

[65] Patton, A. and A. Timmermann (2007). Testing forecast optimality under unknown loss. *Journal of the American Statistical Association* 102(480), 1172–1184.

[66] Phillips, P.C.B., Moon, H.R., 1999. Linear regression limit theory for nonstationary panel data. *Econometrica* 67 (5), 1057–1111.

[67] Philips, P.C.B., Perron, P., 1988, Testing for a unit root in time series regression. *Biometrika*, 75, pp. 335-346.

[68] Pindyck, R.S., Rotemberg, J.J., 1990. The excess co-movement of commodity prices. *Economic Journal* 100 (403), 1173–1189.

[69] Ratti, R.A., Vespignani, J.L., 2016. Oil prices and global factor macroeconomic variables. *Energy Economics* 59, 198-212.

[70] Ravazzolo F., Rothman, P., 2012. Oil and U.S. GDP: A Real-Time Out-of-Sample Examination. *Journal of Money, Credit and Banking* 45(2-3), 449-463.

[71] Stock, J., Watson, M., 2006. Forecasting with many predictors. *Handbook of Economic Forecasting*, vol. 1, 515–554.

[72] Timmermann, A., 2006. Forecast combinations. In: Elliott, C.G.G., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*, vol. 1. Elsevier, pp. 135–196.

[73] Vahid, F., Engle, R.F., 1993. Common trends and common cycles. *Journal of Applied Econometrics* 8 (4), 341–360.

[74] Vahid, F., Engle, R.F., 1997. Codependent cycles. *Journal of Econometrics* 80 (2), 199–221.

[75] Vahid, F., Issler, J.V., 2002. The importance of common cyclical features in var analysis: a monte-carlo study. *Journal of Econometrics* 109 (2), 341–363.

[76] West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica* 64(5), pp. 1067–1084.

[77] Zou, H., 2006. The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101, pp. 1418-1429.

A Online Appendix

A.1 The Derived Demand for Commodities with Zero Short-Run Supply Price Elasticity

It is argued here that there is an important role for demand shocks in explaining the short-run variation of oil and metal prices. These are important globally traded commodities, whose supply function is very price inelastic in the short run – zero or close to zero elasticity. The reason for that lies on the fact that the production function of these commodities is very capital intensive and it takes time to build capital. So, the short-run supply of these commodities, which includes inventories, can be treated as given. Of course, the long-run supply price-elasticity will not be zero or small, so the results in this section will not be verified as the horizon progresses.

To explain this phenomenon, from a theoretical point-of-view, a representative cost-minimizing global industrial firm is considered, which chooses the optimal quantity of inputs x_i, $i = 1, 2, \cdots, n$, all stacked in a vector $x = (x_1, x_2, \cdots, x_n)'$, when producing output y_0. The choice of output y_0 can be thought as an optimal decision coming from the firm's output market. The corresponding prices for inputs $i = 1, 2, \cdots, n$, stacked in a vector $w = (w_1, w_2, \cdots, w_n)'$, are considered given for the firm when choosing x. The firm's cost minimization problem in this context is:

$$\min_{x} \quad C(w, x) = w \cdot x \qquad s.t. \qquad f(x) \geq y_0. \tag{15}$$

From the first-order (interior) condition of this problem, using Shepard's Lemma, the optimal derived demands for all inputs is calculated, labelled $x_i^*(w, y_0)$:

$$\frac{\partial C(w, x^*)}{\partial w_i} = x_i^*(w, y_0), \qquad i = 1, 2, \cdots, n. \tag{16}$$

As argued above, in modelling short-run fluctuations, it is reasonable to assume that oil or metal supply (labelled here as commodity i) cannot be increased without climbing a very steep cost function. Thus, their supply is assumed fixed $\left(\overline{x_i^s}\right)$ in the short run. Thus, the short-run equilibrium condition for inputs (including oil, oil derivatives, and metal commodities) is:

$$x_i^*(w, y_0) = \overline{x_i^s}. \qquad i = 1, 2, \cdots, n. \tag{17}$$

Ceteris paribus, given the equilibrium condition (17), this section investigates how changes in output potentially change the price of input i, w_i. Differentiate (17) considering only changes in w_i and in industrial production, y_0, later solving for $\frac{dw_i}{dy_0}$:

$$\begin{aligned} 0 &= \frac{\partial x_i^*(w, y_0)}{\partial w_i} dw_i + \frac{\partial x_i^*(w, y_0)}{\partial y_0} dy_0, \text{ or,} \\ \frac{dw_i}{dy_0} &= -\frac{\frac{\partial x_i^*(w,y_0)}{\partial y_0}}{\frac{\partial x_i^*(w,y_0)}{\partial w_i}}. \end{aligned} \tag{18}$$

It is straightforward to establish unequivocally that $\frac{dw_i}{dy_0} > 0$ since, from theory, one should have $\frac{\partial x_i^*(w,y_0)}{\partial y_0} > 0$ and $\frac{\partial x_i^*(w,y_0)}{\partial w_i} < 0$. This result $\left(\frac{dw_i}{dy_0} > 0\right)$ is completely intuitive: given concavity of the cost function vis-a-vis input prices $\left(\frac{\partial x_i^*(w,y_0)}{\partial w_i} = \frac{\partial^2 C(w,x^*)}{\partial w_i^2} < 0\right)$, if the representative firm wants to increase industrial production in the short run, it will put an upward pressure on the price of oil, stemming from the fact that it should take more inputs to produce more $\left(\frac{\partial x_i^*(w,y_0)}{\partial y_0} > 0\right)$, otherwise the firm is not a cost minimizer.[8]

To fix ideas, there is below a simple theoretical example, where production follows the well-known Cobb-Douglas production function, with constant returns to scale:

$$y = A x_1^{\alpha} x_2^{1-\alpha}, \qquad 0 < \alpha < 1. \tag{19}$$

The first input (input 1) will represent the input with a short-run fixed supply (oil, oil derivatives, and metal commodities), i.e., the commodity whose production is capital intensive. Input 2 represents an input whose supply responds to price changes in the short run. Indeed, it can be a mix of all other inputs.

The Lagrangian for the cost-minimizing firm reads as:

$$\mathcal{L} = w_1 x_1 + w_2 x_2 - \lambda \left(y_0 - A x_1^{\alpha} x_2^{1-\alpha}\right),$$

[8] The short-run analysis (*ceteris paribus*) omits changes in the price of other inputs. Total differentiation yields:

$$0 = \frac{\partial x_i^*(w, y_0)}{\partial w_i} dw_i + \underbrace{\sum_{j \neq i} \frac{\partial x_i^*(w, y_0)}{\partial w_j} dw_j}_{\text{assumed } = 0 \text{ above}} + \frac{\partial x_i^*(w, y_0)}{\partial y_0} dy_0,$$

But, the effect of $\sum_{j \neq i} \frac{\partial x_j^*(w,y_0)}{\partial w_j} dw_j$ on the price of oil will be of second order: they will have to operate through the cross-effects of substitution and/or complementarity.

where λ is the Lagrange Multiplier. The first-order conditions are:

$$w_1 = \lambda \alpha A \left(\frac{x_2}{x_1}\right)^{1-\alpha}, \tag{20}$$

$$w_2 = \lambda (1-\alpha) A \left(\frac{x_2}{x_1}\right)^{-\alpha}, \tag{21}$$

$$y_0 = A x_1^{\alpha} x_2^{1-\alpha}. \tag{22}$$

Divide (20) by (21) to obtain $x_1 = \frac{w_2}{w_1}\frac{\alpha}{1-\alpha}x_2$. Plugging this last result into (19), and rearranging, gives the optimal demand for input 2:

$$x_2^*(w_1, w_2, y_0) = \frac{y_0}{A} w_1^{\alpha} w_2^{-\alpha} \left(\frac{1-\alpha}{\alpha}\right)^{\alpha}.$$

Symmetrically, one can obtain the optimal demand for input 1:

$$x_1^*(w_1, w_2, y_0) = \frac{y_0}{A} w_2^{1-\alpha} w_1^{\alpha-1} \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha}.$$

Now, if one assumes a short-run fixed supply for input 1 (say, $\overline{x_1^s}$), then, the short-run market equilibrium condition is:

$$\overline{x_1^s} = \frac{y_0}{A} w_2^{1-\alpha} w_1^{\alpha-1} \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha}. \tag{23}$$

Taking logs of (23), solving for $\ln(w_1)$ gives:

$$\ln(w_1) = \frac{1}{1-\alpha}\ln(y_0) - \frac{1}{1-\alpha}\ln\left(\overline{x_1^s}\right) - \frac{1}{1-\alpha}\ln(A) + \ln(w_2) + \ln\left(\frac{\alpha}{1-\alpha}\right),$$

where $\ln\left(\overline{x_1^s}\right)$ and $\ln(A)$ are constants, but, $\ln(w_2)$ depends on $\ln(y_0)$, through an equilibrium condition similar to (23). First, if one disregards the effect of y_0 changes on w_2, there is a positive elasticity $\frac{\partial \ln(w_1)}{\partial \ln(y_0)} = \frac{1}{1-\alpha} > 0$, related to a derived-demand factor, generating synchronicity between commodity prices and global industrial production. However, in principle, one should also consider the effect of changes in $\ln(y_0)$ on $\ln(w_2)$. As long as input 2 supply responds less than proportionally than demand to a change in $\ln(y_0)$, one should have $\frac{\partial \ln(w_2)}{\partial \ln(y_0)} > 0$. Therefore:

$$\begin{aligned}\frac{\partial \ln(w_1)}{\partial \ln(y_0)} &= \frac{1}{1-\alpha} + \frac{\partial \ln(w_1)}{\partial \ln(w_2)}\frac{\partial \ln(w_2)}{\partial \ln(y_0)} \\ &= \frac{1}{1-\alpha} + \frac{\partial \ln(w_2)}{\partial \ln(y_0)} > 0,\end{aligned}$$

and a positive short-run elasticity survives, jointly with the synchronicity of cycles for $\ln(w_1)$ and $\ln(y_0)$.

A.2 Joint Short- and Long-Run Restriction for Oil-Price or for Metal-Price Dynamics using VAR Models

Before discussing the dynamic representation of oil-price or of metal-price data, this section presents useful definitions of the concepts used to measure the degree of comovement in them. These include long-run comovement (cointegration) and short-run comovement (common cycles for the growth rate of these prices). Engle and Issler (1995) and Issler and Vahid (2001) present previous applications of the techniques discussed in this section. For an in-depth theoretical discussion of these issues see Engle and Granger (1987) and Vahid and Engle (1993) respectively for cointegration and common cycles. Theoretical extensions of the standard common-cycle case can be found in Hecq et al. (2006) and Athanasopoulos et al. (2011).

Assume that y_t is a n-vector of $I(1)$ oil price (or log oil price) or of metal prices (or log metal prices), with the stationary $(MA(\infty))$ Wold representation given by:

$$\Delta y_t = C(L)\epsilon_t, \tag{24}$$

where $C(L)$ is a matrix polynomial in the lag operator, L, with $C(0) = I_n$, $\sum_{j=1}^{\infty} \|C_j\| < \infty$. The vector ϵ_t is a $n \times 1$ a multivariate white noise process. One can rewrite equation (24) as:

$$\Delta y_t = C(1)\epsilon_t + \Delta C^*(L)\epsilon_t \tag{25}$$

where $C^*(L) = C_0^* + C_1^* L + C_2^* L^2 + \cdots$, with $C_i^* = \sum_{j>i} -C_j$ for all $i \geq 0$, and, in particular, $C_0^* = I_n - C(1)$.

Integrating both sides of equation (25), given an initial condition y_0:

$$\begin{aligned} y_t &= C(1)\sum_{s=0}^{t-1}\epsilon_{t-s} + C^*(L)\epsilon_t + y_0 \\ &= T_t + C_t \end{aligned} \tag{26}$$

Equation (26) is the multivariate version of the Beveridge-Nelson trend-cycle representation (Beveridge and Nelson, 1981). Apart from an initial condition y_0, the series y_t are represented as sum of a Martingale part $T_t = C(1)\sum_{s=0}^{t-1}\epsilon_{t-s}$, which is called the "trend," and a stationary and ergodic part $C_t = C^*(L)\epsilon_t$, which is called the "cycle."

Definition 1. The variables in y_t are said to have common trends (or cointegrate) if there are r linearly independent vectors, $r < n$, stacked in an $r \times n$ matrix α', with the following property[9]:

$$\underset{r\times n}{\alpha'} C(1) = 0.$$

Definition 2. The variables in y_t are said to have common cycles if there are s linearly independent vectors, $s \leq n - r$, stacked in an $s \times n$ matrix $\tilde{\alpha}'$, with the property that:

$$\underset{s\times n}{\tilde{\alpha}'} C^*(L) = 0.$$

[9]This definition could alternatively be expressed in terms of an $n \times r$ matrix γ, such that:

$$C(1)\gamma = 0.$$

The Granger-Representation Theorem shows that if the series in y_t are cointegrated, α and γ satisfy:

$$\begin{aligned} C(1)\gamma &= 0, \text{and}, \\ \alpha' C(1) &= 0. \end{aligned}$$

Thus, cointegration and common cycles represent restrictions on the elements of $C(1)$ and $C^*(L)$ respectively.

Now, this section discusses what role these restrictions play on the dynamic autoregressive representation of y_t. It is assumed that y_t is generated by a Vector Autoregression (VAR). Note that VARs are the working horses of time-series econometric analysis. They have been applied extensively for reduced-form and structural-form estimation and forecasting, since they fit most macroeconomic and financial data fairly well:

$$y_t = \Gamma_1 y_{t-1} + \ldots + \Gamma_p y_{t-p} + \epsilon_t, \tag{27}$$

where the autoregressive matrix polynomial is $\Phi(L) = I - \Gamma_1 L - \Gamma_2 L^2 - \ldots - \Gamma_p L^p$.

If elements of y_t cointegrate, then the matrix $\Phi(1) = I - \sum_{i=1}^{p} \Gamma_i$ must have less than full rank. In this case, Engle and Granger showed that the system (27) can be written as a Vector Error-Correction model (VECM) as:

$$\Delta y_t = \Gamma_1^* \Delta y_{t-1} + \ \ldots \ + \Gamma_{p-1}^* \Delta y_{t-p+1} + \gamma\alpha' y_{t-1} + \epsilon_t \tag{28}$$

where γ and α are full rank matrices of order $n \times r$, r is the rank of the cointegrating space, $-\left(I - \sum_{i=1}^{p} \Gamma_i\right) = \gamma\alpha'$, and $\Gamma_j^* = -\sum_{i=j+1}^{p} \Gamma_i$, $j = 1, \ldots, p-1$.

For our purposes, testing for cointegration will be used to verify whether oil- or metal-price data share common trends (or have long-run comovement). As is well known, oil or metals are an important input in industrial processes, and thus it is expected that oil or most metal commodities would have their long-run prices linked to global industrial factors. Testing for common trends among y_t will use the maximum-likelihood approach in Johansen (1991).

A key issue to assure that inference is done properly in this case is to estimate the lag length of the VAR (27) consistently, i.e., to estimate p consistently. When data have common cycles as well as common trends, Athanasopoulos et al. (2011) showed that some popular information criteria do not have an appropriate small-sample behavior, and that a combination of traditional information criteria and criteria with data-dependent penalties can estimate the lag length consistently for VARs with common trends and cycles. An alternative way to infer p is to perform diagnostic testing to rule out the risk of underestimation of p, which leads to inconsistent estimates for the parameters in (28).

Vahid and Engle (1993) show that the dynamic representation of y_t may be further restricted if there exist white noise independent linear combinations of the series Δy_t, i.e., that the Δy_t share common cycles. To see this, recall that the cofeature vectors $\tilde{\alpha}_i'$, stacked in an $s \times n$ matrix $\tilde{\alpha}'$, eliminate all serial correlation in Δy_t, i.e. lead to $\tilde{\alpha}'\Delta y_t = \tilde{\alpha}'\epsilon_t$. Therefore, they should restrict the elements of (28) as follows:

$$\tilde{\alpha}'\Gamma_1^* \ = \ \tilde{\alpha}'\Gamma_2^* = \ldots = \tilde{\alpha}'\Gamma_{p-1}^* = 0, \text{ and} \tag{29}$$

$$\tilde{\alpha}'\gamma \ = \ 0. \tag{30}$$

Hecq et al. (2006) have labelled the joint restrictions (29) and (30) as strong-form *serial-correlation common features* (SCCF), whereas they call only imposing restrictions (29) as weak-form SCCF. For the latter, notice that one only inherits an unpredictable linear combination of Δy_t once one *controls* for the long-run deviations $\alpha' y_{t-1}$. Hence,

$$\tilde{\alpha}' \left[\Delta y_t - \gamma\alpha' y_{t-1}\right] = \tilde{\alpha}'\epsilon_t. \tag{31}$$

The discussion of common cycles now continues in the case of strong-form serial-correlation common features, since the weak-form case can be immediately inferred from it. As is well known, cofeature vectors are identified only up to an invertible transformation[10]. Without loss of generality, consider $\tilde{\alpha}$ to have an s dimensional identity sub-matrix:

$$\tilde{\alpha} = \begin{bmatrix} I_s \\ \tilde{\alpha}^*_{(n-s)\times s} \end{bmatrix}$$

Now, $\tilde{\alpha}'\Delta y_t = \tilde{\alpha}'\epsilon_t$ constitute s equations in a system. Completing the system by adding the unconstrained VECM equations for the remaining $n-s$ elements of Δy_t, one obtains,

$$\begin{bmatrix} I_s & \tilde{\alpha}^{*\prime} \\ \underset{(n-s)\times s}{\mathbf{0}} & I_{n-s} \end{bmatrix} \Delta y_t = \begin{bmatrix} & \underset{s\times(np+r)}{\mathbf{0}} & \\ \Gamma_1^{**} & \ldots \; \Gamma_{p-1}^{**} & \gamma^* \end{bmatrix} \begin{bmatrix} \Delta y_{t-1} \\ \vdots \\ \Delta y_{t-p+1} \\ \alpha' y_{t-1} \end{bmatrix} + v_t, \tag{32}$$

where Γ_i^{**} and γ^* represent the partitions of Γ_i^* and γ respectively, corresponding to the bottom $n-s$ reduced form VECM equations, and $v_t = \begin{bmatrix} I_s & \tilde{\alpha}^{*\prime} \\ \underset{(n-s)\times s}{\mathbf{0}} & I_{n-s} \end{bmatrix} \epsilon_t$.

It is easy to show that (32) parsimoniously encompasses (28). Since $\begin{bmatrix} I_s & \tilde{\alpha}^{*\prime} \\ \underset{(n-s)\times s}{\mathbf{0}} & I_{n-s} \end{bmatrix}$ is invertible, it is possible to recover (28) from (32). Notice however that the latter has $s \cdot (np+r) - s \cdot (n-s)$ fewer parameters.

Assuming that y_t share common trends and cycles leading to (32), one tests for common cycles using a canonical-correlation approach once one determines what is the lag length of the VAR, i.e., p. The procedure is as follows:

1. Compute the sample squared canonical correlations between $\{\Delta y_t\}$and $\{\alpha' y_{t-1}, \Delta y_{t-1}, \Delta y_{t-2}, \cdots, \Delta y_{t-p+1}\}$, labelled λ_i, $i = 1, \cdots, n$, where n is the number of variables in the system.

2. Test whether the first smallest s canonical correlations are jointly zero by computing the test statistic:

$$-T \quad \sum_{i=1}^{s} \log\left(1-\lambda_i\right),$$

   which has a limiting χ^2 distribution with $s\,(np+r) - s\,(n-s)$ degrees of freedom under the null, where r is the number of cointegrating relationships. The maximum number of zero canonical correlations that can possibly exist is $n-r$.

3. Suppose that s zero canonical correlations were found in the previous step. Use these s contemporaneous relationships between the first differences as s pseudo-structural equations in a system of simultaneous equations. Augment them with $n-s$ equations from the VECM and estimate the system using full information maximum likelihood (FIML). The restricted VECM will be the reduced form of this pseudo-structural system.

[10] The same is true regarding cointegrating vectors. One is only able to identify a subspace of $\mathbb{R}^n$ of dimension r.

4. In the case where one has weak-form restrictions, in step 1 above, one computes the sample squared canonical correlations between $\left\{\left(\Delta y_t', (\alpha' y_{t-1})'\right)'\right\}$ and $\{\alpha' y_{t-1}, \Delta y_{t-1}, \Delta y_{t-2}, \cdots, \Delta y_{t-p+1}\}$. The rest of the steps remain identical.

A.3 GMM Based Common-Cycle Tests

One possible drawback of the canonical-correlation approach is that it assumes homoskedastic data, and that may not hold for oil-price data or for metal-price data (and other macroeconomic and financial data) collected at high frequency. In this case, a GMM approach is more robust, since inference can be conducted with Heteroskedastic and Auto-correlation (HAC) robust estimates of variance-covariance matrices of parameter estimates. Regarding common cycles, the system with n-equations in (32) can be estimated by GMM. The vector of instruments comprise the series in $\{\alpha' y_{t-1}, \Delta y_{t-1}, \Delta y_{t-2}, \cdots, \Delta y_{t-p+1}\}$ which are collected in a vector Z_{t-1}. Indeed, GMM estimation exploits the orthogonality between all the elements in v_t and all the elements in Z_{t-1}, where v_t is defined in equation (32):

$$\mathbb{E}\left[v_t \otimes Z_{t-1}\right] = \mathbf{0}. \tag{33}$$

The test for common cycles is the over-identifying restriction test – the J test – proposed in Hansen (1982). It has an asymptotic χ^2 distribution with degrees of freedom equal to the number of over-identifying restrictions. As usual, over-identifying restriction tests verify whether or not errors are orthogonal to instruments in an instrumental-variable setup. Thus, it checks whether the exclusions of the elements of Z_{t-1} in the first s equations are appropriate. Heuristically, since the cyclical behavior (serial correlation behavior) of the data Δy_t is captured by Z_{t-1}, this test verifies whether the linear combinations in these s equations have no serial correlation, i.e., are unpredictable. Therefore, it is a test of common serial correlation or *common cycles*. If two series have a *common cycle*, their impulse response functions are colinear, making their response to shocks proportional and therefore similar. Here, contrary to the canonical correlation approach above, one can deal with heteroskedasticity of unknown form by employing the HAC robust estimates for the variance-covariance matrix of sample means counterparts of (32) $\left(\widehat{S_T}\right)$ using the Newey and West (1987) procedure. The parameters estimated by GMM, stacked in a vector θ, comprise all parameters in $\tilde{\alpha}^{*\prime}$ and all parameters in the matrices $\Gamma_1^{**}, \Gamma_2^{**}, \ldots\ \Gamma_{p-1}^{**}$, and in γ^*.

If one wants to test for weak-form SCCF, the only additional twist is that now v_t takes the form:

$$v_t = \begin{bmatrix} I_s & \tilde{\alpha}^{*\prime} \\ \underset{(n-s)\times s}{\mathbf{0}} & I_{n-s} \end{bmatrix} \Delta y_t - \begin{bmatrix} \mathbf{0} & \cdots & \mathbf{0} & \gamma_1 \\ \Gamma_1^{**} & \cdots & \Gamma_{p-1}^{**} & \gamma_2 \end{bmatrix} \begin{bmatrix} \Delta y_{t-1} \\ \vdots \\ \Delta y_{t-p+1} \\ \alpha' y_{t-1} \end{bmatrix},$$

where $\gamma = \begin{bmatrix} \gamma_1 \\ \gamma_2 \end{bmatrix}$.

A.4 Optimal Forecast Combinations

This section lists the set of assumptions needed to obtain optimal forecast combinations and discuss some details on how to construct these optimal forecast under mean-squared error (MSE) risk function. For an in-depth theoretical discussion of these issues see Issler and Lima (2009) and the references therein.

One is interested in forecasting y_t, stationary and ergodic, using information up to h periods prior to t, then, under a MSE risk function, the optimal forecast is the conditional expectation using information available up to $t-h$: $\mathbb{E}_{t-h}(y_t)$. Forecasts of y_t, computed using conditioning sets lagged h periods, are denoted by $f_{i,t}^h$, $i = 1, 2, \ldots, N$. Therefore, $f_{i,t}^h$ are h-step-ahead forecasts and N is the number of models estimated to forecast $f_{i,t}^h$.

Hendry and Clements (2004) argue that the fact that the simple forecast average $\frac{1}{N}\sum_{i=1}^{N} f_{i,t}^h$ usually outperforms individual forecasts $f_{i,t}^h$ shows our inability to approximate $\mathbb{E}_{t-h}(y_t)$ reasonably well with individual models. However, since $\mathbb{E}_{t-h}(y_t)$ is optimal, this is exactly what these individual models should be doing.

With this motivation, our setup writes the $f_{i,t}^h$'s as approximations to the optimal forecast as follows:

$$f_{i,t}^h = \mathbb{E}_{t-h}(y_t) + k_i^h + \varepsilon_{i,t}^h, \tag{34}$$

where k_i^h is the individual model time-invariant bias for h-step-ahead prediction and $\varepsilon_{i,t}^h$ is the individual model error term in approximating $\mathbb{E}_{t-h}(y_t)$, where $\mathbb{E}(\varepsilon_{i,t}^h) = 0$ for all i, t, and h. Here, the optimal forecast is a *common feature* of all individual forecasts and k_i^h and $\varepsilon_{i,t}^h$ arise because of forecast misspecification.

One can always decompose the series y_t into $\mathbb{E}_{t-h}(y_t)$ and an unforecastable component ζ_t, such that $\mathbb{E}_{t-h}(\zeta_t) = 0$ in:

$$y_t = \mathbb{E}_{t-h}(y_t) + \zeta_t. \tag{35}$$

Combining (34) and (35) yields,

$$\begin{aligned} f_{i,t}^h &= y_t - \zeta_t + k_i^h + \varepsilon_{i,t}^h, \qquad \text{or,} \\ f_{i,t}^h &= y_t - \eta_t^h + k_i^h + \varepsilon_{i,t}^h, \;\text{ where, } \eta_t^h = -\zeta_t. \end{aligned} \tag{36}$$

This yields the well known two-way decomposition, or error-component decomposition, of the forecast error $f_{i,t}^h - y_t$:

$$\begin{aligned} f_{i,t}^h &= y_t + \mu_{i,t}^h, \qquad\qquad i = 1, 2, ..., N \text{ and } t > T_1 \\ \mu_{i,t}^h &= k_i^h + \eta_t^h + \varepsilon_{i,t}^h. \end{aligned} \tag{37}$$

By construction, the framework in (36) specifies explicit sources of forecast errors that are found in both y_t and $f_{i,t}^h$; see also the discussion in Palm and Zellner (1992) and Davies and Lahiri (1995). The term k_i^h is the time-invariant forecast bias of model i. It captures the long-run effect of forecast-bias of model i. Its source is $f_{i,t}^h$. The term η_t^h arises because forecasters do not have future information on y_t between $t-h+1$ and t. Hence, the source of η_t^h is y_t , and it is an additive aggregate zero-mean shock affecting equally all forecasts. The term $\varepsilon_{i,t}^h$ captures all the remaining errors affecting forecasts, such as those of idiosyncratic nature and others that affect some but not all the forecasts (a group effect). Its source is $f_{i,t}^h$.

From the perspective of combining forecasts, the components k_i^h, $\varepsilon_{i,t}^h$ and η_t^h play very different roles. If one regards the problem of forecast combination as one aimed at diversifying risk, i.e., a finance approach, then, on the one hand, the risk associated with $\varepsilon_{i,t}^h$ can be diversified, while that associated with η_t^h cannot. On the other hand, in principle, diversifying the risk associated with k_i^h can only be achieved if a bias-correction term is introduced in the forecast combination, which reinforces its usefulness.

The assumptions needed to construct optimal forecast combinations under this framework are now listed.

Assumption 1 One assumes that k_i, $\varepsilon_{i,t}$ and η_t are independent of each other for all i and t.

Independence is an algebraically convenient assumption used throughout the literature on two-way decompositions. At the cost of unnecessary complexity, it could be relaxed to use orthogonal components, something one avoids here.

Assumption 2 k_i is an identically distributed random variable in the cross-sectional dimension, but not necessarily independent, i.e.,

$$k_i^h \sim \text{i.d.}(B, \sigma_{k^h}^2), \tag{38}$$

where B^h and $\sigma_{k^h}^2$ are respectively the mean and variance of k_i^h. In the time-series dimension, k_i^h has no variation, therefore, it is a fixed parameter.

The idea of dependence is consistent with the fact that forecasters learn from each other by meeting, discussing, debating, etc. Through their ongoing interactions, they maintain a current collective understanding of where their target variable is most likely heading to, and of its upside and downside risks. Given the assumption of identical distribution for k_i^h, B^h represents the market (or collective) bias. Since the focus is on combining forecasts, a pure idiosyncratic bias does not matter but a collective bias does. In principle, one could allow for heterogeneity in the distribution of k_i – means and variances to differ across i. However, that will be a problem in testing the hypothesis that forecast combinations are biased.

It is desirable to discuss the nature of the term k_i^h, which is related to the question of why one cannot focus solely on unbiased forecasts, for which $k_i^h = 0$. The role of k_i^h is to capture the long-run effect, in the time dimension, of the bias of econometric models of y_t. The model-based forecasts bias results from model misspecification. Here, it is important to distinguish between in-sample and out-of-sample model fitting. The fact that, in sample, a model approximates well the data-generating process (DGP) of y_t does not guarantee that it will in out-of-sample forecasting; see the discussion in Clements and Hendry (2006) and in Hendry and Clements (2004). Notice that bias correction is a form of intercept correction. Now the discussion is about survey-based forecasts. In this case, a relevant question to ask is: why would forecasters introduce bias under a MSE risk function? Laster et al. (1999), Patton and Timmermann (2007), and Batchelor (2007) list different arguments consistent with forecasters having a non-quadratic loss function. Following their discussion, one assumes that all forecasters employ a combination of quadratic loss and a secondary loss function. Bias is simply a consequence of this secondary loss function and of the intensity in which the forecaster cares for it. The first example is that of a bank selling an investment fund. In this case, the bank's forecast of the fund return may be upward-biased simply because it may use this forecast as a marketing strategy to attract new clients for that fund. Although the bank is penalized by deviating from $\mathbb{E}_{t-h}(y_t)$, it also cares for selling the shares of its fund. The second example introduces bias when there is a market for *pessimism* or *optimism* in forecasting. Forecasters want to be labeled as optimists or pessimists in a "branding" strategy to be experts on "worst-" or on "best-case scenarios," respectively. Batchelor lists governments as examples of experts on the latter.

Assumption 3 The aggregate shock η_t^h is a stationary and ergodic MA process of order at most $h-1$, with zero mean and variance $\sigma^2_{\eta^h} < \infty$.

Since h is a bounded constant in our setup, η_t^h is the result of a cumulation of shocks to y_t that occurred between $t-h+1$ and t. Being an $MA(\cdot)$ is a consequence of the wold representation for y_t. If y_t is already an $MA(\cdot)$ process, of order smaller than $h-1$, then, its order will be the same of that of η_t^h. Otherwise, the order is $h-1$. In any case, it must be stressed that η_t^h is unpredictable, i.e., that $\mathbb{E}_{t-h}\left(\eta_t^h\right) = 0$.

Assumption 4: Let $\varepsilon_t^h = \left(\varepsilon_{1,t}^h, \varepsilon_{2,t}^h, \ldots\, \varepsilon_{N,t}^h\right)'$ be a $N \times 1$ vector stacking the errors $\varepsilon_{i,t}^h$ associated with all possible forecasts, where $\mathbb{E}\left(\varepsilon_{i,t}^h\right) = 0$ for all i and t. Then, the vector process $\left\{\boldsymbol{\varepsilon}_t^h\right\}$ is assumed to be covariance-stationary and ergodic for the first and second moments, uniformly on N. Further, defining as $\xi_{i,t}^h = \varepsilon_{i,t}^h - \mathbb{E}_{t-1}\left(\varepsilon_{i,t}^h\right)$, the innovation of $\varepsilon_{i,t}^h$, one assumes that

$$\lim_{N\to\infty} \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} \left|\mathbb{E}\left(\xi_{i,t}^h \xi_{j,t}^h\right)\right| = 0. \tag{39}$$

Because the forecasts are computed h-steps ahead, forecast errors $\varepsilon_{i,t}$ can be serially correlated. Assuming that $\varepsilon_{i,t}^h$ is weakly stationary is a way of controlling its time-series dependence. It does not rule out errors displaying conditional heteroskedasticity, since the latter can coexist with the assumption of weak stationarity; see Engle (1982).

The techniques discussed in this section are appropriate for forecasting a weakly stationary and ergodic univariate process $\{y_t\}$ using a large number of forecasts that will be combined to yield an optimal forecast in the mean-squared error (MSE) sense. These forecasts are the result of several econometric models that need to be estimated prior to forecasting[11].

Three consecutive distinct time sub-periods are considered, where time is indexed by $t = 1, 2, \ldots, T_1, \ldots, T_2, \ldots, T$. The first sub-period E is labeled the "estimation sample", where models are usually fitted to forecast y_t subsequently. The number of observations in it is $E = T_1 = \kappa_1 \cdot T$, comprising $(t = 1, 2, \ldots, T_1)$. For the other two, one follows the standard notation in West (1996). The sub-period R (for regression) is labeled the post-model-estimation or "training sample", where realizations of y_t are usually confronted with forecasts produced in the estimation sample, and weights and bias-correction terms are estimated. It has $R = T_2 - T_1 = \kappa_2 \cdot T$ observations in it, comprising $(t = T_1 + 1, \ldots, T_2)$. The final sub-period is P (for prediction), where genuine out- of-sample forecast is entertained. It has $P = T - T_2 = \kappa_3 \cdot T$ observations in it, comprising $(t = T_2 + 1, \ldots, T)$. Notice that $0 < \kappa_1, \kappa_2, \kappa_3 < 1$, $\kappa_1 + \kappa_2 + \kappa_3 = 1$, and that the number of observations in these three sub-periods keep a fixed proportion with T - respectively, κ_1, κ_2 and κ_3 - being all $O(T)$. This is an important ingredient in our asymptotic results for $T \to \infty$.

Issler and Lima (2009) propose a non-parametric estimator of k_i^h, which exploits the fact that k_i^h represents the fixed effect of a panel of forecasts:

$$\left(f_{i,t}^h - y_t\right) = k_i^h + \eta_t^h + \varepsilon_{i,t}^h, \qquad \begin{aligned} &i = 1, 2, \ldots, N \\ &t = T_1 + 1, \ldots, T_2 \end{aligned} \tag{40}$$

[11] In this setting, one can also imagine that some (or all) responses use no formal econometric model at all, e.g., just the result of an opinion poll on the variable in question using a large number of individual responses.

It does not depend on any distributional assumption on $k_i^h \sim i.d.(B^h, \sigma^2_{k^h})$ and it does not depend on any knowledge of the models used to compute the forecasts $f_{i,t}^h - y_t$.

Issler and Lima (2009) propose the following consistent estimators for the components k_i^h, B^h, η_t^h, and $\varepsilon_{i,t}^h$:

$$\widehat{k}_i^h = \frac{1}{R}\sum_{t=T_1+2}^{T_2} f_{i,t}^h - \frac{1}{R}\sum_{t=T_1+2}^{T_2} y_t \tag{41}$$

$$\widehat{B^h} = \frac{1}{N}\sum_{i=1}^{N} \widehat{k}_i^h \tag{42}$$

$$\widehat{\eta}_t^h = \frac{1}{N}\sum_{i=1}^{N} f_{i,t}^h - \widehat{B^h} - y_t \tag{43}$$

$$\widehat{\varepsilon}_{i,t}^h = f_{i,t}^h - y_t - \widehat{k}_i^h - \widehat{\eta}_t^h \tag{44}$$

Now the two most important results from Issler and Lima (2009) are stated, which asserts that the (feasible) bias-corrected average forecast (BCAF) is an optimal forecasting device.

Theorem 1. *If Assumptions 1-4 hold, the feasible bias-corrected average forecast* $\sum_{i=1}^{N} \omega_i f_{i,t}^h - \widehat{B^h}$ *obeys*

$$\underset{(T,N\to\infty)_{seq}}{plim} \left(\frac{1}{N}\sum_{i=1}^{N} f_{i,t}^h - \widehat{B^h} \right) = y_t + \eta_t^h = \mathbb{E}(y_t)$$

and has a mean-squared error as follows:

$$\mathbb{E}\left[\underset{(T,N\to\infty)_{seq}}{plim} \left(\frac{1}{N}\sum_{i=1}^{N} f_{i,t}^h - \widehat{B^h} \right) - y_t \right]^2 = \sigma^2_{\eta^h}$$

Therefore it is an optimal forecasting device.

Indeed, there are infinite ways of combining forecasts. The next corollary presents alternative weighting schemes.

Corollary 1. *Consider the sequence of deterministic weights* $\{\omega_i\}_{i=1}^N$,*such that* $\omega_i \neq 0$, $\omega_i = O(N^{-1})$ *uniformly, with* $\sum_{i=1}^N \omega_i = 1$ *and* $\lim_{N\to\infty} \sum_{i=1}^N \omega_i = 1$. *Then, under Assumptions 1-4, an optimal forecasting device is:*

$$\mathbb{E}\left[\underset{(T,N\to\infty)_{seq}}{plim} \left(\sum_{i=1}^{N} \omega_i f_{i,t}^h - \sum_{i=1}^{N} \omega_i \widehat{k}_i^h \right) \right]^2 = \sigma^2_{\eta^h}$$

Optimal population weights, constructed from the variance-covariance structure of models with stationary data, will obey the structure in Corollary 1 and cannot perform better than $\frac{1}{N}$ coupled with bias correction. Therefore, there is no forecast-combination puzzle in the context of populational weights.

Theorem 1 shows that the feasible BCAF is asymptotically equivalent to the optimal weighted forecast. Its advantage is that it employs equal weights. As $N \to \infty$, the number of estimated parameters is kept at unity: $\widehat{B^h}$. This is a very attractive feature of the BCAF compared to devices that combine forecasts using estimated weights.

Our answer to the curse of dimensionality is parsimony, implied by estimating only one parameter- $\widehat{B^h}$. One additional advantage is that one needs not limit the asymptotic path of N and T, which is the case of forecasts based on estimated weights.

Finally, there is one interesting case in which one can dispense with estimation in combining forecasts: when the mean bias is zero, i.e., $B^h = 0$, there is no need to estimate B^h and the BCAF is simply equal to $\frac{1}{N}\sum_{i=1}^{N} f_{i,t}^h$, the sample average of all forecasts. This is the ultimate level of parsimony. To be able to test the null that $B^h = 0$, Issler and Lima developed a robust t-ratio test that takes into account the cross-sectional dependence in k_i^h.

To assess forecast accuracy, an algorithm appropriate for the bias-corrected average forecast (BCAF) is constructed as follows:

1. For each model, estimate the coefficients of the regressors using the sub-sample from 1 to T_1.
2. Forecast h-steps ahead the models estimated in step 1 (f_{it}^h) from T_1 to T_2. Each model should be forecasted h-steps ahead $T_2 - T_1 - h + 1$ times.
3. Calculate the bias associated with each h-step ahead forecasts and each model; the bias is the average error between the h-steps ahead forecast and the observed value of the target series (from T_1 to T_2).
4. Forecast h-steps ahead the same models estimated in step 1 for only T_2+h, using the same coefficients estimated in step 1.
5. Store the bias from step 3 and the forecast made in step 4, f_{i,T_2+h}^h.
6. Update $T_1 = T_1 + 1$, $T_2 = T_2 + 1$.
7. Go to step 1 until $T_2 = T$.
8. Adjust the forecasts of each model (made from $T_2 + 1$ to T) by their respective bias.
9. Combine all these adjusted forecasts using equal weights.
10. Compute the RMSE of the BCAF, considering the series of instantaneous price variation (first difference of logarithms) of Brent Crude Oil as the target series.

Forecast Combinations for Nested Models

It is important to discuss whether the techniques above are applicable to the situation where some (or all) of the models combined are nested. The potential problem is that the innovations from nested models can exhibit high cross-sectional dependence. In what follows, nested models are introduced into our framework in the following way. Consider a continuous set of models and split the total number of models N into M classes (or blocks), each of them containing m nested models, so that $N = mM$. In the index of forecasts, $i = 1, \ldots, N$, nested models are grouped contiguously. Hence, models within each class are nested but models across classes are non-nested. The number of classes and the number of models within each class are set to be functions of N, respectively as follows: $M = N^{1-d}$ and $m = N^d$, where $0 \leq d \leq 1$. Notice that this setup considers all the relevant cases: (i) $d = 0$ corresponds to the case in which all models are non-nested; (ii) $d = 1$ corresponds to the case in which all models are nested and; (iii) the intermediate case $0 < d < 1$ gives rise to N^{1-d} blocks of nested models, all with size N^d.

For each block of nested models, Assumption 4 may not hold because the innovations from that block can exhibit high cross-sectional dependence. Regarding the interaction across blocks of nested models, it is natural to impose that the correlation structure of innovations across classes is such that Assumption 4 holds, since one should expect that the cross-sectional dependence of forecast errors across classes is weak.

Here, keeping some nested models poses no problem at all, since the mixture of models will still deliver the optimal forecast. From a practical point of view, the choice of $0 \leq d < 1$ seems to be superior. Here, one is back to the main theorem in finance about risk diversification: do not put all your eggs in the same basket, choosing a large enough number of diversified (classes of) models.

A.5 MIDAS Regressions

This section discusses mixed-frequency regression methods for predicting economic activity. Daily financial data as well as monthly macroeconomic data is used here to estimate the mixed data sampling (MIDAS) regression model. The objective is to highlight the value of daily information and build real-time forecasts of the current (and future) real GDP growth rates.

A number of articles have documented the advantages of using MIDAS regressions in terms of improving quarterly macro forecasts with monthly data, or improving quarterly and monthly macroeconomic predictions with a small set of daily financial series; see, for instance, Clements and Galvão (2008); Hamilton (2008); Ghysels and Wright (2009) and Kuzin et al. (2011).

Following the notation of Andreou et al. (2013), assume the variable of interest is observed at some low frequency (e.g., real GDP quarterly growth rate), denoted by Y_t^Q, and the goal is to forecast this variable $h-$quarters ahead, that is Y_{t+h}^Q. To do so, the econometrician has daily observations of the financial predictors X. Denote by $X_{m-j,t}^D$ the jth day counting backward in quarter t, where m denotes the number of trading days per quarter – assumed to be constant for the sake of simplicity (e.g., $m = 66$). Hence, $X_{m,t}^D$, considering $j = 0$, corresponds to the observation of X on the last day of quarter t.

The ADL-MIDAS (p_Y^Q, q_X^D) model is given by:

$$Y_{t+h}^Q = \mu^h + \sum_{j=0}^{p_Y^Q-1} \rho_{j+1}^h Y_{t-j}^Q + \beta^h \sum_{j=0}^{q_X^D-1} \sum_{i=0}^{m-1} \omega_{i+j*m}^{\theta^h} X_{m-i,t-j}^D + \varepsilon_{t+h}^h, \tag{45}$$

which entails p_Y^Q lags of Y_t^Q and q_X^D lags of $X_{m-i,t}^D$. The weighting scheme ω involves a low-dimensional vector of unknown parameters θ, used to avoid the parameter proliferation implied by the estimation of coefficients associated to high frequency lags. There are many possible parameterizations of the model. Following Andreou et al. (2013), the exponential Almon lag polynomial is adopted, since this parameterization yields a parsimonious, yet flexible scheme of data-driven weights. The parameters of the ADL-MIDAS model in equation (45) are estimated by nonlinear least squares.

First, note that this model can be used to obtain direct (as opposed to iterated) forecasts for h multiperiods ahead. Also, the ADL-MIDAS generalizes the ADL forecasting approach to deal with mixed-frequency data. See Bai et al. (2013) for a good discussion about how the MIDAS regression model relates to the traditional Kalman filter mixed-frequency approach.

The MIDAS model described in equation (45) can be extended to include macroeconomic data available at monthly frequency. Moreover, as new releases of data become available, one can include useful information between periods t and $t+1$ (e.g., due to nonsynchronized publication lags). In this sense, Clements and Galvão (2008) introduced the MIDAS regression with *leads*, where the notion of *leads* pertains to the fact that one can use information between t and $t+1$. For instance, suppose one is 2 months into quarter $t+1$ (i.e., at the end of February, May, August or November) and the goal is to forecast quarterly GDP. This implies one has approximately 44 trading days (2 months) of daily data.

The ADL-MIDAS (p_Y^Q, q_X^D, J_X^M, J_X^D) regression model with *leads* (in both monthly and daily data) is described as follows (see Andreou et al., 2013, equation 5.1):

$$\begin{aligned} Y_{t+h}^{Q} &= \mu^{h} + \sum_{j=0}^{p_Y^Q-1} \rho_{j+1}^{h} Y_{t-j}^{Q} + \sum_{j=3-J_X^M}^{2} \gamma_j^h X_{3-j,t+1}^{M} \\ &\quad + \beta^{h} \left[\sum_{i=(3-J_X^D)*m/3}^{m-1} \omega_{i-m}^{\theta^h} X_{m-i,t+1}^{D} + \sum_{j=0}^{q_X^D-1}\sum_{i=0}^{m-1} \omega_{i+j*m}^{\theta^h} X_{m-i,t-j}^{D} \right] + \varepsilon_{t+h}^{h}, \end{aligned} \tag{46}$$

where $X_{m-i,t+1}^D$ denotes the ith day counting backward in quarter $t+1$, and J_X^D represents daily *leads*, for the daily predictor, in terms of multiples of months (i.e., $J_X^D = 1$ and 2). In other words, in the case of $J_X^D = 2$, $X_{2m/3,t+1}^D$ corresponds to $2m/3 = 44$ leads (assuming $m = 66$), while $X_{1,t+1}^D$ corresponds to 1 lead for the daily predictor. The definitions of J_X^M and $X_{3-j,t+1}^M$, associated to monthly data, are quite similar to the ones used for daily data.

This empirical exercise forecasts the real GDP growth rate (Y_{t+h}^Q), for the U.S. or the OECD, $h-$quarters ahead, using the MIDAS approach. To do so, one uses no lags of the real GDP growth ($\rho_{j+1}^h = 0$) and consider as monthly variable $X_{month=i}^M = \Delta\ln(\text{global industrial production index}_{month=i})$, and as daily variables $X_{day=j}^D = \frac{1}{15}\sum_{k=1}^{15} \Delta\ln(\text{commodity price}_{day=j-k+1})$, where the set of commodities is either $[oil]'$ or $[aluminium; copper; tin]'$. The source of daily data is Thomson Reuters,[12] whereas the monthly global industrial production series is extracted from the IMF-IFS database. Regarding model (45), the predictors $[X^D; X^M]'$ are considered lagged two months in respect to the end of the quarter of the observed GDP (Y_t^Q). In model (46), the predictors $[X^D; X^M]'$ are considered one and two months forwarded, respectively, in respect to the end of the quarter of the observed GDP (Y_t^Q). In other words, for model (46), one takes into account information between periods t and $t+1$ when considering *leads* of one and two months, respectively, for the commodity prices and the global industrial production index.

[12] Nominal price data is deflated using the producer price index (PPI) for the U.S., extracted from the FRED database of the St. Louis FED.