EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Elena Maria Diaz, Gabriel Perez Quiros

Daily tracker of global economic activity: a close-up of the COVID-19 pandemic

No 2505 / December 2020

Abstract

This paper develops a novel indicator of global economic activity, the GEA Tracker, which is based on commodity prices selected recursively through a genetic algorithm. The GEA Tracker allows for daily real-time knowledge of international business conditions using a minimum amount of information. We find that the GEA Tracker outperforms its competitors in forecasting stock returns, especially in emerging markets, and in predicting standard indicators of international business conditions. We show that an investor would have inexorably profited from using the forecasts provided by the GEA Tracker to weight a portfolio. Finally, the GEA Tracker allows us to present the daily evolution of global economic activity during the COVID-19 pandemic.

JEL Classification: F44, G17, Q02

Keywords: Global Economic Activity, Commodity Prices, Factor Models, Variable Selection, Genetic Algorithm, Leading Indicators

Non-Technical Summary

In the context of the COVID-19 global pandemic, the need for real-time high-frequency indicators to understand the evolution of the economy has become paramount. Particularly, the heightened uncertainty and continuously changing economic and social conditions, along with the global nature of this crisis, requires the development of such an indicator for global economic activity. This is particularly useful for policy-makers, economic agents and the general public in order to perform optimal decisions that contribute to a faster recovery. In this paper, we present the GEA Tracker, a real-time indicator for global economic activity in daily frequency that uses only a minimum amount of information which is publicly available.

Existing measures of global business conditions contain a number of caveats. Some of them present serious publication delays, and most of them are biased to capture the activity in advanced economies, where the information is more reliable. This under-representation of emerging economies hampers the accuracy and performance of these indicators. A strand of the literature overcomes these issues by focusing on commodity prices, which are timely and publicly available and mostly reflect expectations of firms for future production. Nonetheless, these indicators are very noisy, some of them clearly do not reflect movements in economic activity and more dangerously, their idiosyncratic component is cross-correlated, all of which makes the selection of commodities to build such an indicator a central issue.

In this paper, we present a novel procedure to select the relevant set of commodity price series. We consider a total of 47 commodities, including metals, energy, raw materials and agricultural products. A genetic algorithm is implemented to objectively select the best combination of commodities to proxy global business conditions.

We find that a global economic activity indicator based on commodity prices is significantly more accurate when these are carefully selected. Mainly, a simple combination of selected metals and energy prices provides the best proxy for global business conditions, while the inclusion of all commodities undercuts the accuracy of the indicator. This is particularly due to the incorporation of agricultural products, which are highly volatile and cross-correlated, with a majority of changes unrelated to the global economy.

We build the GEA Tracker using only information that would have been available at each point in time. We find that our algorithm can properly adapt to a shift in the structure of the economy or in technology by changing the pool of selected commodities. With this real-time point of view, the GEA Tracker proves to be highly accurate in signaling periods of recession or expansion in the global economy.

Additionally, the GEA Tracker has desirable forecasting properties. It significantly improves the predictability of global business conditions with respect to benchmark models, particularly for the manufacturing sector. Moreover, we find significant predictability of stock returns of emerging markets, and show that an investor would have inexorably profited from using the forecasts provided by the GEA Tracker to weight a portfolio.

Finally, we show daily real-time estimates of the GEA Tracker which provides a close-up of the evolution of global economic activity during the COVID-19 pandemic.

1. Introduction

After the declaration by the World Health Organization (WHO) on March 11th, 2020 of the COVID-19 as a global pandemic, governments began implementing a series of mobility restrictions with the purpose of slowing down the contagion. The harshness of these restrictions brought global economic activity to its largest slowdown since the Great Depression of 1929. As the first wave of the pandemic passed, how deep the economy would fall or how the recovery would look like became a main point of interest. Today, as the second wave of COVID-19 hits the world and the distribution of a vaccine remains months away, there is a prevailing degree of uncertainty and continuously changing conditions. There exists, hence, a paramount need for closely monitoring the evolution of global economic activity which would allow policy-makers and economic agents to perform the decisions that enable a faster recovery. With this need in mind, in this paper we develop the GEA Tracker, a high-frequency indicator of global economic activity that is publicly available and estimated in real time.

Existing measurements of global economic activity vary greatly within the literature. The most natural approach has been to use world gross domestic product (GDP) or industrial production (e.g. Manescu and Van Robays, 2016; Herrera and Rangaraju, 2019). Also, e.g., Cooper and Priestley (2013) estimate the global business cycle using a capital-to-output ratio. However, these measures contain several caveats that will be discussed later, a main one being that they are published with a significant delay which make them invalid to address world economic conditions in real time.

A different approach has been to extract the information present in the fluctuations of commodity markets. This approach is interesting because global economic activity determines the aggregate demand for all industrial commodities and, therefore, this demand can be used as a proxy for global economic activity. For example, Ravazzolo and Vespignani (2017), put forward the use of world steel production as a monthly indicator for global real economic activity. This measure, however, is subject to idiosyncratic shocks in the steel industry for which a different proposal is to aggregate information on numerous commodity markets. This is done for the Kilian (2009) index, a widely used indicator in the literature that is based on ocean freight rates for the transportation of industrial commodities. However, the raw data used to construct this indicator is not publicly available and it also suffers from publication delays.

Alternatively, one can simply observe the co-movement of commodity prices, using a factor model. Delle Chiaie et al. (2017) and Alquist et al. (2019) show that the common factor of commodity prices serves as a proxy for global economic activity. This has the advantage that commodity prices consist in readily and publicly available data. Moreover, commodity demand reflects the expectations of firms for future production, which causes it to lead global business conditions. This should prove particularly useful not only in signaling but also in forecasting global economic activity.

Nevertheless, Kilian and Zhou (2018) state that the choice of the commodity prices that are to be included in the estimation of a global economic activity indicator is central. In fact, Alquist et al. (2019) discuss the issue of the selection of commodities when extracting a common factor to measure their aggregate demand. To solve this, they provide a set of theoretical criteria for a commodity to be included in the estimation. Yet, as Kilian and Zhou (2018) state, “how important imposing each of these constraints is in practice, remains an open question at this point”. We provide an answer to this question by developing a novel indicator of global economic activity where the set of commodity price series is carefully selected.

In this sense, our paper is in line with Boivin and Ng (2006) who raise the concern that “more is not always better”. This is motivated by the fact that when using more series to extract factors, these become less useful for forecasting. This is shown to be true in practice when the idiosyncratic errors do not satisfy the weak cross correlation condition. When this condition does not hold, the idiosyncratic co-movements dominate the signal, making the factor biased. It might therefore be the case that the common factor extracted from all commodity prices assigns significantly more weight to, e.g., agricultural commodities, which do not accurately indicate fluctuations in global economic activity, than it assigns to energy and metals, which do. We show that this is the case. There is, thus, a concern of including irrelevant commodity price series in the estimation of a factor model to identify global commodity demand.

The optimal selection of commodities, though, poses a highly nonlinear optimization problem with a large number of possible solutions. Bai and Ng (2008) address this issue by examining the use of well-known methodologies for selection and shrinkage such as LASSO, and Elastic Net, which are special cases of the Least Angle Regression (LARS). However, as will be shown, these methodologies, while improving the accuracy of a global economic activity indicator estimated through the selection of commodities, are limited by the fact that they search through the solution space of possible commodity sets in a highly restrictive manner. We resolve this by alternatively implementing a genetic algorithm which efficiently explores all the possible sets of commodity prices to find the one that generates the most accurate measure of global economic activity. This selection is performed in real time with the information available up to the end of each year.

We begin by estimating the GEA Tracker in a monthly frequency from 1960M1 to 2020M8. This is done to provide a sufficiently long time series for the examination of the properties of the indicator, including its ability to identify global business cycles and the accuracy of the genetic algorithm of selecting only those commodities that are truly relevant to the global economic structure.

Followingly, we test the validity of the GEA Tracker by using it to forecast the Global Composite Purchasing Managers’ Index (PMI) and the Global Manufacturing PMI, as leading indicators of international business conditions. We find that the GEA Tracker significantly reduces the out-of-sample mean squared error of an autoregressive (AR) model and outperforms alternative measures of global economic activity.

We then also perform an out-of-sample forecasting exercise for world, developed, and emerging stock markets.[1] We find significant stock return predictability in emerging markets

[1] There is evidence that financial markets incorporate changing global business conditions. (Bekaert and Harvey, 1995, Imbs, 2006, Ang and Bekaert, 2007, Bekaert et al., 2011, Cooper and Priestley, 2013)

when using the GEA Tracker, which reflects these countries' dependency on commodities. We also show that a mean-variance investor would have profited most from using our forecasts to weight a portfolio.

Finally, once the properties of the GEA Tracker have been carefully examined, we proceed to estimating it at a daily frequency from January 2nd, 2020 to September 15th, 2020. This is done with the purpose of closely observing the evolution of global economic activity during the COVID-19 global pandemic. We show that the evolution of the GEA Tracker closely follows the reading of the main economic shocks provoked by the global pandemic.

The structure of the paper is as follows. Section 2 describes the estimation of the GEA Tracker at a monthly frequency. Section 3 contains the estimation results and properties of the GEA Tracker. Section 4 presents its forecasting properties. Section 5 describes the estimation of the GEA Tracker at a daily frequency for year 2020, providing a close-up of the evolution of global economic activity during the pandemic. Finally, section 6 concludes.

2. Estimation of a Global Economic Activity Tracker

2.1. The Model

Global economic activity is estimated through the following factor model:

$$p_{it} = \lambda_i f_t + e_{it} \quad \forall\, i \,\exists\, \{1, \ldots, n_s\} \tag{1}$$

where p_{it} is the log-level of the real price of commodity i at time t, λ_i is a loading factor, f_t is global economic activity, e_{it} is the idiosyncratic component for commodity i, and n_s is the total number of selected commodities. Equation 1 states that, at a certain period, the change in price of a commodity results from either a shift in the common factor f_t or from supply or demand shocks that pertain exclusively to the commodity. We estimate this model using Principal Component Analysis (PCA). Prices are deflated using the U.S. consumer price index (CPI), as proxy for World CPI.

Note that Equation 1 may have two important drawbacks. The first one is that we assume that there are no dynamics in the evolution of f_t or e_{it}, which would allow to forecast the evolution of global economic activity several periods ahead in time. Later in the text, in section 3.4, we relax this assumption by introducing dynamics to both the common factor f_t and the idiosyncratic components e_{it}, where a dynamic factor model (DFM) is estimated using the Kalman filter rather than PCA. We will see that the gains in the estimation are minor compared to the gains obtained from carefully analyzing the second drawback: the selection of variables.

Certainly, the estimation of a DFM and the selection of commodities are not mutually exclusive. However, attempting both approaches for such a large data set becomes too computationally intensive, basically because of the difference in estimation time between PCA and a fully dynamic factor model using the Kalman filter. We, therefore, focus on the selection of commodities. Once the best model is selected, we introduce dynamics in section 3.4 to show the robustness of the results and to be able to forecast, if necessary, global economic activity several periods ahead.

2.2. Selection of Commodity Prices

For the selection of commodities, we begin by choosing a benchmark global economic activity indicator as reference for the aggregate demand of commodities, defined as f_t in Equation 1. We then set the genetic algorithm to look for the combination of commodities whose co-movement best matches the fluctuations of this indicator.

Some macroeconomists lean towards the use of proxies for world real gross domestic product (GDP) or measures of world industrial production (see, e.g, Manescu and Van Robays, 2016; Herrera and Rangaraju, 2019). There are, however, many caveats with these measures, as discussed in Kilian and Zhou (2018). For instance, in the case of real GDP, the percent contribution of the service sector, which has no direct relation with the demand for commodities, has steadily risen since the 1970s. Also, for the case of industrial production, its link with global commodity markets may be weakened by changes in the stock of raw materials held by governments and firms. In addition, the time series for these two indicators are too short, with sample data for global real GDP starting in 1990Q1, and for the OECD + BRICS industrial production indicator starting in 2006. Moreover, both real GDP and industrial production constitute coincident indicators to the global business cycle, and therefore, by construction, will be lagging the timing of shifts in the aggregate demand of commodities, which should lead business conditions. They are also not available in real-time. More importantly, neither provide appropriate weights to emerging countries, or they make use of inadequate data for these economies.

An alternative well-established global economic activity indicator is the Kilian (2009) index. This indicator is constructed using information on ocean freight rates for industrial commodities, based on a long observed positive correlation between economic activity and shipping rates. Not only is it constructed as far back as January 1968 but, unlike global real GDP and world industrial production, it automatically accounts for shifting country weights, changes in the composition of real output, and changes in technology and productivity. More importantly, the Kilian index is a coincident indicator with respect to the volume of shipping in global commodity markets and responds instantaneously to shifts in the aggregate demand of commodities (Kilian and Zhou, 2018). Hence, it reflects expectations of firms for future production, making it a leading indicator with respect to global output. The magnitude of the fluctuations of commodity demand implied by this index are therefore more accurate than those indicated by global GDP or industrial production.

A concern with the Kilian index is that the shipbuilding cycle accentuates upswings in the index caused by commodity booms and exaggerates downswings caused by commodity busts (Kilian and Zhou, 2018). However, in practice, the business cycle timings implied by the Kilian index are consistent with historical anecdotal evidence of global expansions and recessions. Concerns that the sustained decline of the Kilian index after 2010 may be due to excess capacity in global dry cargo shipping markets are discussed by Kilian and Zhou (2018), who show that this reflects a genuine economic slowdown due to a sluggish growth of the Chinese economy. We therefore set the Kilian index as the benchmark indicator to reference the selection of commodity prices. Consequently, our novel indicator approximates these business cycle timings while not being subject to changes in the shipbuilding cycle, because it is constructed with commodity price data. Moreover, since the GEA Tracker incorporates information both on commodity prices as well as ocean freight rates, it is therefore built as the aggregate demand of the set of selected commodities that also reflects the shifts in the demand of the maritime transportation of industrial commodities.

It is important to note, however, that we do not directly use the Kilian index for tracking global economic activity given two reasons. One is that the primary data source for the construction of this indicator is found in Drewry's Shipping Consultants Ltd. which is not publicly available. The other is that it only becomes available in the author's personal website with a publication delay; whereas, in the current context as well as in forecasting, readily-available data is crucial. One approach is to extrapolate the Kilian indicator using the Baltic Dry Index (BDI) as in Baumeister and Kilian (2014), which is a shipping and trade index estimated by the Baltic Exchange that measures the changes in the cost of maritime transportation of several raw materials, for given shipping paths, time of delivery, and speed. We will later show that while this approach does improve the forecasting performance when attempting to predict global business or financial conditions, it continues to be outperformed by the GEA Tracker. This responds to the fact that the GEA Tracker is based on commodity prices which constitute data that is readily-available and publicly accessible.

2.2.a. The Genetic Algorithm

The selection of commodities is performed through a genetic algorithm, which was first made popular by Holland (1975). It is a method used for solving constrained and unconstrained optimization problems based on biological evolution. Its population-based search technique mimics the principle of natural selection laid by Charles Darwin.

The basic idea of this algorithm is that, in the same way a living organism is the result of the combination of its genetic information, we can think of a possible GEA Tracker as the result of the combination of the information it contains, which is the set of commodity price series used to construct such indicator. Evolutionary theory states that, given a population, the most fit individuals have the highest probability of reproducing and passing on their genetic information to future generations. This is why, with time, a certain population maintains and propagates strong genes (or combination of genes), while weaker genes eventually die out. Similarly, the genetic algorithm we design works by generating a population of possible GEA Trackers, each containing different genetic information (different combinations of commodity price series used in its estimation). Each GEA Tracker is then evaluated against the benchmark indicator to determine its fitness. More fit indicators are given a higher probability of passing on their genetic information to future generations. New populations are then iteratively created. Holland (1975) showed that genetic algorithms could be applied to highly complex and nonlinear optimization problems such as ours. He proved that after enough iterations, the genetic information will converge to the optimal solution.

Here, we provide a brief overview of the genetic algorithm for the selection of commodities. A detailed explanation can be found in Appendix A.

In our application, we have a total of $n = 47$ commodity price series. We therefore have a data set of size $T \times n$, where T is the sample size. We define an individual (or potential solution) A_j as a genome of n binary genes a_{ij}, which can take the value of 0 or 1, $\forall\, i \ni \{1, \ldots, n\}$.

$$A_j = (a_{1j}, a_{2j}, ..., a_{nj}) \tag{2}$$

where a_{ij} is assigned the value of 1 if commodity i is included in the estimation of the indicator, and 0, if it is not. This implies that for individual A_j the original data set of size $T \times n$ is reduced by eliminating all columns i, where $a_{ij} = 0$. The resulting data set is then used to estimate Equation 1, where the common factor f_t is identified as f_{jt} for individual A_j.[2]

The objective is then to find the optimal individual A^* that maximizes the R^2 statistic when performing the following regression:

$$f_t^* = \mu + \beta f_{jt} + \epsilon_t \tag{3}$$

where f_t^* is the Kilian index.

We start the algorithm by generating a random initial population where each individual is assigned a fitness value defined as the R^2 resulting from Equation 3. Once all fitness values are estimated, these are turned into probabilities through an operation known as scaling. This allows for individuals to be stochastically chosen through a selection operator, to be part of the reproduction process, where a new set of individuals, called children, will be created. The most fit individuals have a higher probability of being selected, and those who are selected are called parents.

The reproduction process consists in performing two operations: crossover and mutation. Crossover reproduction combines the genetic information from a pair of parents, generating two new individuals from each pair. This operation allows the genetic algorithm to explore the search space of possible solutions. Mutation, on the other hand, is applied to a single parent by altering its genes with a small probability. This is done to provide diversity to the population and avoid premature convergence to a local solution.

Additionally, we set our genetic algorithm to be elitist, which means that a small group of elite individuals are ensured survival to the next generation. This guarantees that the most fit members are maintained throughout the evolution process.

Finally, the individuals resulting from crossover, mutation and elitism substitute the existing population, forming a new generation. The entire process is then repeated by estimating the fitness values of the new individuals and again applying the scaling, selection,

[2]Note that because each gene has two possible values, we have a set of potential solutions of size 2^n individuals. This totals to 1.47×10^{14} combinations. The evaluation of such a large set, nevertheless, is made computationally viable by the genetic algorithm.

crossover, mutation and elitism operators. This is done throughout several generations until the population converges to an optimal solution or a stopping criterion is met.

2.2.b. LASSO: An Alternative Method for the Selection of Commodities

One could consider alternative methods for the selection of commodities. LASSO, for example, is widely used in the literature for selection and shrinkage. Particularly, in Bai and Ng (2008), the authors evaluate the performance of LASSO, Elastic Net and Least Angle Regression for the selection of series in a factor model. As the authors state, both LASSO and Elastic Net are merely special cases of Least Angle Regression, and we will therefore, for simplicity, only consider LASSO.

Accordingly, we also estimate Equation 1 by selecting the pool of commodities using LASSO selection.

LASSO provides a ranking of the series according to its predictive power for the targeted variable. Principal components is then estimated using the first three series in the ranking to obtain the first potential global economic activity indicator. The pool of commodities is expanded to include the following series in the ranking, generating a new possible indicator in each iteration. Then, as in Bai and Ng (2008), we use the BIC criteria, resulting from equation 3, to select the factor.

In sections 3 and 4, the properties of the indicator estimated using LASSO will be compared to that of the GEA Tracker.

2.2.c. Data Set of Commodity Prices

We use monthly data for commodity prices spanned between January 1960 and August 2020, available from the World Bank. We only consider those commodities for which data is available for the full sample.[3] However, because the total of 47 commodities implies a rather large number of possible solutions, we assist the genetic algorithm in finding the optimal combination by first reducing it to a subset of commodities which contains only metals and energy, and then extending it until the full set of commodities is considered. We do this because it is reasonable to expect that metals and energy, as they constitute main inputs intensively used in industrial production processes, will be more highly correlated with global economic activity than other industrial commodities (e.g., fertilizers, raw materials, fats and oils) or than agricultural commodities.

The data set of the World Bank contains the following types of commodities: energy, beverages, fats and oils, grains, other foods, raw materials, fertilizers, metals and minerals, and precious metals. For our application, we define the following subsets:

$$\Omega_M \subseteq \Omega_I \subseteq \Omega \tag{4}$$

where Ω_M includes the commodities that we denote as Metals & Energy and is a subset of all industrial commodities, Ω_I. Then, both Ω_M and Ω_I are subsets of Ω, which includes all

[3]For the case of crude oil, we include the price series estimated by the World Bank as the average between Brent crude, WTI crude and Dubai crude.

47 commodities. This classification can be seen in further detail in Figure 1.

< Insert Figure 1 about here >

The genetic algorithm is set to first select the optimal combination of commodities from Ω_M which contains a total of 13 price series (Figure 1). Once the genetic algorithm has selected the optimal combination of commodities within subset Ω_M (which we denote A^*_M), we run it again for subset Ω_I, but by first seeding the initial population with individual A^*_M. This is done by setting one of the individuals of the initial population equal to A^*_M. The remaining initial population is generated randomly from a uniform distribution, as described earlier. Note that, because our genetic algorithm is elitist, it will search the solution space for a combination of industrial commodities that performs better than A^*_M; otherwise, the selected combination A^*_I will be set to $A^*_I = A^*_M$. Finally, we run the genetic algorithm for the full set Ω, in order to find the optimal combination of all commodities, A^*, for which we seed the initial population with A^*_I.[4]

With this procedure, we generate three GEA Trackers through the selection of commodities among the sets of metals and energy (Ω_M), industrial commodities (Ω_I) and all commodities (Ω). These are denoted GEA Tracker Metals & Energy, GEA Tracker Industrial and GEA Tracker, respectively. For completeness, we keep all three indicators for evaluation and comparison. This will allow us to determine if there are any significant gains from including industrial commodities, other than metals and energy, or agricultural commodities in the estimation of a global economic activity indicator. We also add a fourth indicator, for comparison, which is generated by estimating the model described in Equation 1, but with no selection at all; that is, where $a_{ij} = 1 \; \forall \; i \; \exists \; \{1, ..., n\}$.

2.3. Recursive Estimation of the Indicator

It is important to note, that to ensure that, for the spanned sample, on would have been able to construct this indicator in real time, we use only information available at each point in time to not only assign the weights to commodity prices in the estimation of the GEA Tracker, but to perform the selection as well.

This implies that, at any time t, our GEA Tracker, f_t, is constructed with the last subset of commodities that has been selected by the genetic algorithm as the optimal combination, given by A^*_t. Note that A^*_t is generated by the genetic algorithm using a set of information available at time t, which we denote as I_t. In this sense f_t can be defined as a function of $(A^*_t \mid I_t)$. However, because there is data available for all commodity price series before time t, one might be tempted to use this same combination of commodities A^*_t to estimate the GEA Tracker at any time $t-h$. However, h periods before, we did not have the information, I_t, that allowed us to determine the optimal combination of commodities, A^*_t. If that were the case, we would define $f_{t-h}(A^*_t \mid I_t)$, where the GEA Tracker at time $t-h$ would be

[4] Also, at any time t we define our set Ω to include all commodities that have a positive historical correlation with our benchmark indicator y_t. This is done in order to include only commodities that are pro-cyclical. Such is the case for all metals and energy.

estimated with information that only becomes available in the future, h periods ahead. To avoid this, we perform the selection of commodities recursively and, therefore, create the series $\{f_1(A_1^* \mid I_1), \ldots, f_T(A_T^* \mid I_T)\}$.

We do this by beginning the exercise in 1979M12. Later, every year, we update the selection, A_t^*, from December of 1980 to December of 2019.[5] Also, because we consider a three month publication delay of the Kilian index, the last three months of data are extrapolated by regressing the benchmark indicator on the BDI, once it becomes available in 1985. For all previous years, the genetic algorithm is implemented by obtaining the fitness values from the regression of the Kilian index over the estimated indicators up to September of each year. The pool of selected commodities is then kept during each following year and PCA is used, for the following twelve months, to estimate the weights given to each commodity price series at a monthly frequency.

3. Estimation Results and Properties of the GEA Tracker

3.1. Explained Variance of the Benchmark Indicator

Figure 2 shows all estimated indicators along with the benchmark Kilian index, where the value of each indicator at any time t, $f_t(A_t^* \mid I_t)$, corresponds to the most recent estimate, for each case. Note that even though the pool of commodities selected for estimation changes over time, the value of the indicator remains comparable from one period to the next because data is standardized in order to perform PCA.

< Insert Figure 2 about here >

As can be seen in Figure 2, the indicators estimated with the selection of commodities present fluctuations that match those of the Kilian index more closely than an indicator estimated using all commodity price series. However, in order to carefully measure the gains of variable selection, we estimate the variance explained of the Kilian index at each period t, by each of the five proposed global economic activity indicators. The nature of this comparison is the following: in each period t, the selection algorithm produces an optimal combination of variables, which, through principal components, results in a time series from the first period up to period t, $f(A_t^* \mid I_t)$. This proposed 'optimal' vintage in each period of time is then used as an explanatory variable for Kilian's index. Figure 3 then shows the explained variance of the benchmark indicator by each proposed indicator at time t.

< Insert Figure 3 about here >

As can be seen in Figure 3, the selection of commodities in the estimation of a global economic activity indicator largely increases the explained variance of the Kilian index with

[5]The selection of commodities could also be performed every month. However, the estimation would become too computationally intensive. The basket of commodities is then selected only once a year for parsimony.

respect to an indicator generated with no selection at all. The GEA Tracker with selection using the genetic algorithm explains around 50% of the variance of the Kilian index, while the model with no selection barely explains around 10% of the variance on average. Results then support that there is a gain in selecting variables prior to applying principal components, which rebuts the idea of "more is always better" that has been standard in many factor-model analyses in the literature.

Particularly, one can even observe a gain in explained variance when estimating the global economic activity indicator through the selection of variables using LASSO. This gain, however is not nearly as high as that obtained using the genetic algorithm. Results show, that with a smaller sample, using no selection even outperforms the selection through LASSO. As the sample increases, nonetheless, LASSO selection shows a learning ability, and therefore begins to outperform the indicator generated with no selection. This provides further proof of the importance of variable selection. Nevertheless, the selection of commodities using the genetic algorithm is clearly more adequate.

Additionally, one can observe that the gain that is obtained from extending the pool of commodities to include raw industrial materials, fertilizer and agricultural products, even when performing selection with the genetic algorithm, is minimum. This is particularly true starting 2005, when the commodity boom started, and remains true through to the end of the sample. This provides evidence that although agricultural products may have helped in signaling fluctuations in global economic activity in the first part of the sample, it is no longer true for the last part. Also, although for the years previous to the commodity boom, the genetic algorithm incorporates agricultural products in the estimation of the indicator, this only marginally contributes to achieving a higher explained variance of the benchmark indicator. Most explained variance, in the entire sample, is then due to the estimation of a common factor between selected metals and energy.

Furthermore, because the selection of commodities is done once a year and kept every month within that year, a sharp change in the price of one of the selected commodities significantly alters the values of the indicator. This generates a special concern over agricultural products which have high variances and, therefore, large price shifts that are mainly idiosyncratic. This explains the drop in performance that is observed in the early and late 1980s and mid 1990s of the indicators estimated by the selection within Ω_I and Ω below the one estimated through the selection of only metals and energy, Ω_M. As the sample enlarges, and the genetic algorithm therefore counts on more information, the selected commodities within the full set Ω correspond mainly to those selected amongst metals and energy.

An interesting observation is the drop in correlation of all estimated indicators with the benchmark index in 2005. At this moment, global economic activity was characterized by the commodity boom, an increase in the price of physical commodities mainly driven by the rising demand from emerging markets such as the BRIICS[6] countries, particularly China and India. This commodity boom started in January 2001 and continued until June 2008, when commodity prices crashed due to the Great Financial Crisis. This phenomenon is signaled by

[6] Brazil, Russia, India, Indonesia, China and South Africa.

the Kilian index and all proposed indicators with an upward trend during this period (Figure 2). However, the Kilian index also signals a relative drop in global economic activity from December 2004 to August 2005, which does not coincide with anecdotal evidence at the time.[7] The reason for this can be found in events in the ocean freight markets that year. As reported by UNCTAD (2006), although the seaborne shipments of the main bulks, particularly iron ore and coal, increased by 7.2 percent in 2005; the balance between supply and demand for both time and trip charters resulted in lower freight rates. In fact, the drop on the average annual index of rates of time and trip charters reported was of 20.0 and 12.2 percent, respectively. The downward turn in 2005 then reflected in the Kilian index reduces the correlation with commodity-price based indicators, which continue to estimate an upward trend in 2005 as they are not affected by events in the ship-building cycle.

Furthermore, the decrease in correlation between our GEA Tracker and the benchmark index after 2011 is explained by unprofitable levels of freight rates for ship owners from 2011 to 2015. At this time, ship owners had invested in large capacity ships which, as global economic conditions weakened due to the slowdown in Chinese growth, generated a vessel oversupply and increased volatility in ocean freight rates (see UNCTAD, 2016). The ship-building cycle then caused the Kilian index to magnify the drop in global economic activity and introduce higher volatility to its estimations. Alternatively, while this drop is also evidenced by commodity price behavior and identified by our GEA Tracker, it is done with a much smoother and less volatile downward trend.

3.2. Business Cycle Identification

We now examine the accuracy with which the GEA Tracker identifies global business cycles. Figure 4 shows the NBER dated recessions for the United States along with the Kilian index and the GEA Tracker, Metals & Energy. As we can observe, they both coincide with anecdotal evidence concerning expansions and recessions and provide evidence of a slowdown in global economic activity in the last few years of the sample.

< Insert Figure 4 about here >

Both the Kilian index and the GEA Tracker effectively capture the second recession of the 1970s, caused by the oil embargo of the OPEC and the fall of the Bretton Woods system; the crisis of the early 2000s, which has been attributed to the dot-com bubble; and the great financial crisis in 2008. They also both provide evidence of the slowdown from 2011 to 2016 due to sluggish Chinese growth.

Also note that, for the first recession in the 1970s, the GEA Tracker better coincides with the NBER dates. Additionally, in the double dip recession of the 80s, the Kilian index correctly captures the second dip of the economy by signaling a downward trend since the start of 1981. However, the GEA Tracker also signals a drop in global economic activity since February of 1980, effectively capturing the downturn during both periods. More importantly, there is a more visible downward trend by the GEA Tracker during the recession in 1990.

[7] World output grew by 3.6 percent in 2005.

This is because the ratio of freight rates to most commodity prices increased in 1990 with respect to 1989 (UNCTAD, 1990). This was reported to be due to high shipping insurance costs and anti-competitive shipping practices.

Finally, we do not observe any false signals of recessions given by the GEA Tracker. Any significant downturns seem to coincide with known recessions or slowdowns of the global economy. Overall, the GEA Tracker removes some of the contamination in the Kilian index caused by events in the supply side of the freight rate industry and provides a robust measure of the global business cycle.

3.3. Selection of Commodities Relevant to the Global Economy

It is important to note here that the genetic algorithm prompts us to select only those commodities that are truly relevant for global economic activity, diminishing to the minimum the amount of noise in the system. To observe this, we estimate the aggregated weights given by each indicator and at each point in time to every commodity type, defined in the following way:

$$\lambda_M = \sum_i \lambda_i \quad \forall\, i \ni \Omega_M \tag{5}$$

$$\lambda_I = \sum_i \lambda_i \quad \forall\, i \ni (\Omega_I - \Omega_M) \tag{6}$$

$$\lambda_A = \sum_i \lambda_i \quad \forall\, i \ni (\Omega - \Omega_I) \tag{7}$$

where λ_M, λ_I, and λ_A are the aggregated weights assigned to metals and energy, other industrial commodities (e.g., fertilizers, raw materials, fats and oils) or agricultural commodities, respectively, and λ_i is defined as in Equation 1.

For comparison, we examine the respective weights assigned when there is no selection and f_t is merely obtained by estimating Equation 1 through PCA on all commodity price series; when the selection is performed using LASSO, as well as when the genetic algorithm is implemented to perform the careful selection of commodity prices to estimate Equation 1. Figure 5 shows the results.

< Insert Figure 5 about here >

We can observe in panel A that when all commodities are included in estimating a global economic activity indicator through principal components with no selection, information on agricultural products, raw materials, fertilizers, fats and oils dominates in the factor model. Moreover, the system does not learn over time that agricultural products might be noisy and continues to add signals that are misleading with respect to the evolution of global economic activity. However, in panels B and C we can observe how both LASSO as well as the genetic algorithm learn over time with the input of more information. They both decrease the importance of agricultural products while giving more weight to metals and energy. This becomes even more clear during the commodity boom. Nonetheless, it is clear that the genetic algorithm not only has a significantly higher learning ability, but proves to be more stable and less noisy.

Additionally, we examine the weights, λ_i, given to each individual commodity i selected by the genetic algorithm from the full set, Ω.[8] This is shown in Figure 6. We only show the commodities that are included in the estimation of the indicator at any point of the sample. All other commodities are always given a weight of zero.

< Insert Figure 6 about here >

Figure 6 shows that raw materials and agricultural commodities are selected by the genetic algorithm only sporadically, and none are kept for long periods of time or given large weights. The fact remains that the common factor amongst selected metal and energy commodities best signals fluctuations in global economic activity.

One particular commodity that stands out is platinum. The price series of this commodity is selected by the genetic algorithm throughout the entire sample, and it is given an even larger weight in the estimation of the indicator starting 2005. Platinum is a precious metal with unique physical and catalytic properties and is used primarily by manufacturers of the automobile, electronics and jewelry industries. For the automobile sector, it is central in reducing vehicle carbon dioxide emissions while in electronics it increases the media storage capacity of laptops and servers. Furthermore, platinum catalysts increase yields in chemical processes, increasing efficiency in industrial production. Additionally, it also has applications in biomedical technology and optics.

Other metals that prove noteworthy are copper and zinc, which are given a large weight in the years before the commodity boom. Copper is a very ductile and malleable metal and is a very good conductor of electricity. It is also relatively inexpensive. However, there has been a gradual substitution of copper for aluminum. Whereas major applications of copper are electrical wire, roofing and plumbing, and industrial machinery; aluminum has been mostly used in transportation, packaging, and household items. However, by mid 2000s, the relative price ratio between copper and aluminum began to increase, causing a faster rate of substitution between the commodities, as technological advances were made to allow for further usage of aluminum. Developments in aluminum wiring that compensate for lower conductivity and flexibility and allow for more efficient and less corrosive conductors, have driven major power companies to make the switch from copper to aluminum. There is now more scope to replace copper in power grid cables, auto wiring, air conditioning and refrigeration systems. This has also had its effect in the automotive industry, as aluminum replaces copper as a conductor in on-board power systems. By ends of the sample, the genetic algorithm therefore no longer incorporates the price of copper in the estimation of the GEA Tracker and begins to incorporate the price of aluminum.

[8] We also have the results of the weights given to commodities when selection is performed solely in sets Ω_M and Ω_I. However, we only include the results from set Ω for simplicity. Weights given to commodities are similar in all three cases. Full results are available upon request.

In the case of zinc, one of its most important qualities is its natural capacity to protect steel from corrosion. Zinc coatings provide a physical barrier and cathodic protection to the underlying steel. 60 percent of all zinc is destined to steel protection, which is then used for construction, infrastructure, automobiles and machinery. However, the price of zinc is not only driven by global economic activity, but also by changes in inventories of firms, governments and investors. Indeed, when a U.S. stockpile of zinc for national defense was authorized for disposal in 2005, demand for zinc fell significantly both in the U.S. and Europe. Nickel, on the other hand, which also constitutes one of the main inputs in the production of steel, gained great interest in light of the commodity boom. In fact, this commodity is considered of great economic and strategic importance to many countries, as it is used in a wide variety of industries: mobile phones, food preparation equipment, transport, buildings, power generation, etc. Nickel based alloys are also used for gas turbines, chemical plants, electronics, coinage and marine engineering. Our genetic algorithm, then substituted the price of zinc for the price of nickel in 2005, all in all keeping at least one major input for steel production in the estimation of the GEA Tracker.

Furthermore, natural gas has also become an important commodity in the two past decades. It is extensively used to heat homes and generate electricity, and has commercial and industrial applications, while emitting far less carbon dioxide than fossil fuels. Growing concern for environmental issues and technological advancements in natural gas production and distribution, makes it an important commodity to include in the estimation of the GEA Tracker as an indicator of energy consumption. This is done by the genetic algorithm starting in 2006.

On this note, one might find interesting that crude oil is not given any significant weight in estimating global economic activity. The use of the price of crude oil would have been problematic due to large idiosyncratic shocks from the beginning of the sample up to early 1990s where OPEC and political disturbances greatly determined the price of oil. All in all, our selection procedure proves to successfully discriminate among commodity prices that truly signal global economic fluctuations at any point in time.

3.4. Selection versus Dynamics

So far, we have used a static factor model, described in Equation 1, to estimate the GEA Tracker, without taking into account any dynamics in the factor. Because global economic activity can be assumed to have a smooth behavior, a DFM has been considered in the literature for constructing such an indicator. Delle Chiaie et al. (2017), e.g., estimate a DFM with a block structure to identify a global factor, block-specific components related to specific commodity markets, and purely idiosyncratic shocks. They do not, however, consider variable selection.

Therefore, to estimate the gains of carefully analyzing the variable selection with respect to the gains of estimating a fully dynamic factor model, we relax the assumption of no dynamics on the common factor and idiosyncratic shocks. To do so, we extend the model in Equation 1 with the following two equations:

$$f_t = \varphi_1 f_{t-1} + \cdots + \varphi_p f_{t-p} + \omega_t \qquad \omega_t \sim iid\ N(0, \sigma^2_\omega) \tag{8}$$

$$e_t = \psi_1 e_{t-1} + \cdots + \psi_p e_{t-p} + \varepsilon_t \qquad \varepsilon_t \sim iid\ N(0, \sigma^2_\varepsilon) \tag{9}$$

which describe the dynamics of the common factor f_t and the idiosyncratic terms e_{it}, respectively.

However, the large number of parameters implied by a DFM requires the use of an algorithm, such as the expectation maximization (EM) algorithm proposed by Doz et al. (2011). Note that within the genetic algorithm, this requires for the EM algorithm to be implemented for every single individual to obtain its fitness value, which would have been too computationally intensive.

We therefore simply estimate the GEA Tracker using a DFM for the full sample, in which we use the last selected pool of commodities (natural gas, nickel and platinum). We also estimate a DFM for the case of no selection. We then compare the indicators resulting from the following estimations: (1) static factor model with selection, (2) dynamic factor model with selection, (3) static factor model with no selection, (4) dynamic factor model with no selection. Figure 7 shows the resulting indicators.

< Insert Figure 7 about here >

As can be observed in Figure 7, there is little difference between indicators estimated through a static factor model and a dynamic factor model. There are, however, large differences between the indicators that perform the selection of variables as opposed to those that do not. Altogether, results show that the consideration of dynamics in the factor model is not nearly as important as the selection of variables in the estimation of a global economic activity indicator.

We would also expect no large differences in the resulting pools of commodities selected by a genetic algorithm at each point in time if we had considered a DFM from the start.

4. Forecasting Properties of the GEA Tracker

We now examine the forecasting properties of the GEA Tracker on economic and financial variables.

A highly desirable property of a global economic activity indicator is its ability to forecast global business conditions. If the GEA Tracker is a good indicator of activity, it should be useful for forecasters to predict changes in world business conditions. Therefore, to test the forecasting performance of the GEA Tracker, we select the global PMI index as a target variable. Not only is the PMI one of the most reliable leading indicators for assessing the state of the economy but, because it is based on the views of purchasing managers, it provides information on current and future business conditions.

Additionally, with respect to financial markets, we examine the ability of the GEA Tracker to forecast international stock returns, which is the most used variable to describe global financial conditions.

4.1. Forecasting the Purchasing Managers' Index

The PMI is an indicator of economic health for manufacturing and service sectors. The purpose of the index is to provide information about current business conditions to firms and analysts. It is released monthly by the Institute for Supply Management and is based in five major survey areas (new orders, inventory levels, production, supplier deliveries and employment) with questions on whether business conditions will be improving, deteriorating or remaining equal. We use the estimated global economic activity indicators to forecast the Global Composite PMI and the Global Manufacturing PMI.

One caveat in performing this exercise is that the Global Composite PMI and the Global Manufacturing PMI are only available starting July 1998 and January 1999, respectively, providing a rather short sample for the forecasting exercise. We solve this by extending both indices with the United States PMI, which is available since January 1960. We refer to these extended time series as the Extended Global PMI and the Extended Global Manufacturing PMI.

The exercise is performed for the sample spanned from 2003M1 to 2020M8. To simulate the situation of a real forecaster, only data available before 2003M1 is used for an in-sample estimation of parameters, which is then updated recursively each period as new information is obtained for an out-of-sample estimation. Numerous researchers have stressed the importance of out-of-sample forecasting (Rapach et al., 2010; among others), particularly to avoid the data mining that might result from in-sample estimations.

We use the following predictive regression to generate the forecasts:

$$\Delta y_{t+1} = \alpha + \beta_1 \Delta y_t + \beta_2 \Delta f_t + \xi_{t+1} \qquad \xi_t \sim N(0, \sigma^2) \tag{10}$$

where Δy_{t+1} is the predicted change in the PMI index for time $t+1$, α is a constant, Δf_t is the change in the global economic activity indicator at time t, and ξ_{t+1} is an error term that belongs to a normal distribution with zero mean and variance σ^2. The forecasting performance of this model is then compared to a benchmark AR(1) model, where the forecast for the change in the PMI is given by

$$\Delta y_{t+1} = \alpha + \beta_1 \Delta y_t + \xi_{t+1} \qquad \xi_t \sim N(0, \sigma^2) \tag{11}$$

For global economic activity, f_t, we consider the following indicators: the Kilian index, the common factor of all commodity prices estimated with no selection, the indicator estimated using LASSO selection, and the three GEA Trackers, recursively estimated with a genetic algorithm to select commodities among Ω_M, Ω_I, and Ω. For the Kilian index we consider a 3-month publication delay. We also consider the approach proposed by Baumeister and Kilian (2014) to update the Kilian index using the BDI, and refer to this as the updated

Kilian index.

The forecasting performance of each model is evaluated by the out-of-sample R^2 statistic, R^2_{OS}, suggested by Campbell and Thompson (2008) to compare the $\widehat{\Delta y}_{t+1}$ and $\overline{\Delta y}_{t+1}$ forecasts at a 1-month horizon, where $\widehat{\Delta y}_{t+1}$ is a forecast based on the predictive model described by Equation 10 and $\overline{\Delta y}_{t+1}$ is a forecast based on the benchmark model described by Equation 11. Followingly, the R^2_{OS} statistic is given by

$$R^2_{OS} = 1 - \frac{\sum_{t=t_{o-1}}^{T-1} (\Delta y_{t+1} - \widehat{\Delta y}_{t+1})^2}{\sum_{t=t_{o-1}}^{T-1} (\Delta y_{t+1} - \overline{\Delta y}_{t+1})^2} \tag{12}$$

where Δy_{t+1} is the realized change in the PMI index, t_o is the start of the forecasting sample (2003M1) and T is the end of the sample (2020M8). Consequently, the R^2_{OS} measures the reduction in Mean Squared Prediction Error (MSPE) for the predictive regression model relative to the benchmark model, in percentage terms. Thus, notice that when $R^2_{OS} > 0$, the $\widehat{\Delta y}_{t+1}$ forecasts outperforms the benchmark forecasts $\overline{\Delta y}_{t+1}$.

We then examine whether the results are statistically significant. To do so, we test against the null of equal MSPE between the two models. However, note that, because the benchmark model is nested in the predictive regression model, the parameter β_2 in Equation 10 would be zero in the population under the null. This produces an upward bias in the estimation of the MSPE of the predictive regression model produced by such parameter. Therefore, we estimate the MSPE-adjusted statistic developed by Clark and West (2007). This statistic adjusts for the bias by deducting the square difference in the point predictions generated by each model as follows.

We first define

$$\zeta_{t+1} = (y_{t+1} - \overline{y}_{t+1})^2 - [(y_{t+1} - \widehat{y}_{t+1})^2 - (\widehat{y}_{t+1} - \overline{y}_{t+1})^2] \tag{13}$$

and then regress $\{\zeta_{s+1}\}_{s=t_0}^{T-1}$ on a constant and calculate the t-statistic under the null that the constant is zero. The p-value is obtained with a standard normal distribution, for a one-tailed test. Table 1 reports the estimated R^2_{OS} using each global economic activity indicator, and its statistical significance.

< Insert Table 1 about here >

Results for the PMI indices show that using the GEA Tracker to forecast business conditions significantly outperforms the benchmark model, indicating that we can effectively predict a percentage of the fluctuations in these indices. Note that for the sample period, the performance of the three GEA Trackers is almost equal. This is because most of the weight is given to metal and energy commodities for all cases, for which they have a similar forecasting performance.

Moreover, we significantly improve our results when predicting the manufacturing rather than the composite PMI, agreeing with the fact that we are measuring the demand for industrial commodities which relate to the manufacturing sector of the economy, and not the service sector. We are able to improve up to 10.850% of the out-of-sample forecast accuracy for the Extended Global Manufacturing PMI with respect to the benchmark model, as opposed to only 3.879% that we are able to improve for the Global Composite PMI, which also includes services.

An interesting observation is that the GEA Tracker performs significantly better when extending the PMI for a longer time series as shown by the results for the extended PMI indices, as opposed to the case of the benchmark Kilian index. Because the target time series is extended with data from the U.S. PMI, this may be reflecting a mismatch with the global PMI for which we consider results for the non-extended series more accurate. In this context, it is also important to note that the fluctuations in global business conditions as signaled by the Kilian index are far more accurate than those signaled by the common factor of all commodity prices, when there is no selection. It is only by selecting the commodities which best match the fluctuations in the Kilian index, as we do for our proposed indicator, that we are able to significantly improve forecasting performance.

Results also highlight the importance of a global economic activity indicator being available in real-time. When adopting the approach of Baumeister and Kilian (2014) of updating the Kilian index with the BDI, we improve the forecasting performance for the Global Manufacturing PMI, with respect to using the Kilian index with a publication delay. However, this gain is much shorter in comparison to the gain obtained with our approach of using carefully selected commodity price data.

Finally, in all cases, the indicators generated with selection, either through LASSO or through the genetic algorithm, outperform one based on no commodity selection both in the magnitude of the R^2_{OS}, as well as in its statistical significance. This shows that the inclusion of commodities that have no relation with global business conditions will deteriorate the forecasting performance of an estimated commodity factor. The careful selection of commodities through a genetic algorithm then proves highly relevant in practice. Nonetheless, the GEA Tracker outperforms the indicator generated through LASSO in all cases, once more proving that the genetic algorithm is best suited for variable selection in our application.

4.2. Forecasting Stock Returns

We now perform an out-of-sample forecasting exercise for the sample spanned from 2000M1 to 2020M8 where our objective variable is the real stock return for the following regions: World, Developed Countries, and Emerging Countries. Stock indices are obtained from the MSCI, and real stock returns are measured by deflating the stock index with the U.S. CPI, as proxy for world inflation, and then obtaining the percentage change. All data available before 2000M1 is used for an in-sample estimation of parameters, which is then updated recursively each period as new information is obtained.

We use the following predictive model:

$$\Delta y_{t+1} = \alpha + \beta \Delta f_t + \xi_{t+1} \qquad \xi_t \sim N(0, \sigma^2_{t+1}) \tag{14}$$

$$\sigma^2_{t+1} = \omega + \phi_1 \sigma^2_t + \phi_2 \xi^2_t \tag{15}$$

where Δy_{t+1} is the forecast for the real stock return at time t, α is a constant, Δf_t is the change in the global economic activity indicator at time t, and ξ_t is an error term that belongs to a normal distribution with zero mean and a time-varying volatility characterized by a GARCH(1,1) process, which was developed by Engle (1982) and is typically used by financial institutions to model the volatility of stock returns. For global economic activity, f_t, we consider the same indicators as in the forecasting exercise for the PMI. The forecasting performance of this model is then compared to a random walk model, as benchmark, that is defined as follows:

$$\Delta y_{t+1} = \alpha + \xi_{t+1} \qquad \xi_t \sim N(0, \sigma^2_{t+1}) \tag{16}$$

where volatility, σ^2_{t+1}, is modelled as in Equation 15.

We also evaluate the forecasting performance of each model through the R^2_{OS} statistic described in Equation 12, where now $\widehat{\Delta y}_{t+1}$ is a forecast based on the predictive model described by Equation 14, $\overline{\Delta y}_{t+1}$ is a forecast based on the benchmark model described by Equation 16, t_o is 2000M1 and T is 2020M8. We also test for the statistical significance of the R^2_{OS} through the MSPE-adjusted statistic.

Additionally, we measure the economic importance of using each global economic activity indicator to forecast international stock returns. To do so, we consider that, for any market, an individual can perform investment decisions based on the returns forecasted by a given indicator. The value of the information depends on the certainty of the expected returns, which is the result of the accuracy of the forecasts generated by each model. To determine this value, we construct a series of portfolios based on such forecasts and then estimate the certainty equivalent of each portfolio, which is the risk-free return that the individual would consider equivalent to taking the risk of investing in the portfolio, given its expected returns and volatility.

For this, we consider that, at any point in time, the individual is faced with the decision of investing on the stock market or on a risk-free rate, usually based on government bonds. Given the forecast for the stock return, the uncertainty of that forecast, and the known risk-free rate, the investor then decides how much to allocate on each asset. The return and volatility of the portfolio will depend on the weight the investor gives to the stock market and to the risk-free asset. However, to determine the true value of the information provided by the forecasts, these need to be used optimally to construct the portfolio in a way that maximizes the returns while reducing the volatility. We therefore estimate the optimal weights as in Markowitz (1952), in the following way:

$$w_t = \frac{\widehat{\Delta y}_{t+1} - r_{f,t+1}}{\gamma \widehat{\sigma^2_{t+1}}} \tag{17}$$

where $\widehat{\Delta y}_{t+1}$ is the forecasted stock return for period t, $r_{f,t+1}$ is the risk-free rate known at time t, γ is the risk aversion coefficient which is set to 2, and $\widehat{\sigma^2_{t+1}}$ is the forecasted volatility of the stock return. Note that this weight is typically contained between 0 and 1 but can take negative values if the investor shorts the market or values larger than 1 if he leverages the market.

At the end of each period, the investor's return is then given by the weighted average between the realized stock return and the risk-free rate, as described in the following equation:

$$r_{p,t} = w_{t-1}\Delta y_t + (1 - w_{t-1})r_{f,t} \tag{18}$$

where $r_{p,t}$ is the return of the portfolio at time t, w_{t-1} is the weight assigned by the investor to the stock index in the previous period, and Δy_t is the realized stock return.

Finally, the certainty equivalent of a given portfolio is estimated by

$$ce = \overline{r_p} - \frac{\gamma}{2}\sigma^2(r_p) \tag{19}$$

where ce is the certainty equivalent of the portfolio, and $\overline{r_p}$ and $\sigma^2(r_p)$ are the mean and variance of the returns of the portfolio. Note that, because a typical investor is often risk-averse, the value of a portfolio is given by its expected returns minus a cost of uncertainty.

The certainty equivalents are estimated for portfolios constructed with the point forecasts generated by each global economic activity indicator through the predictive regression model described in Equations 14 and 15. We also estimate the certainty equivalent of the portfolio based on the benchmark random walk model. This is done for the sample spanned between 2000M1 and 2020M8 and for each stock market: World, Developed and Emerging.

We report the certainty equivalent gain of each portfolio as the difference with respect to the certainty equivalent of the random walk portfolio. This is given in percentage values and annualized. For robustness, we also estimate the certainty equivalent gain by restricting the weights, w_t, between -0.5 and 1.5. This is to prevent the investor from choosing unrealistic leverages of the portfolio when longing or shorting the stock market at any point in time.

Results are reported in Table 2.

< Insert Table 2 about here >

Table 2 shows that we can only find significant predictability of monthly stock returns in emerging markets. This is explained by the fact that emerging economies are not only less financially diversified but also far more dependent on commodities. Using the GEA Tracker to forecast emerging market stock returns outperforms the random walk model, by reducing the out-of-sample mean squared error up to 2.368%

Also, as in the PMI forecasting exercise, results highlight the importance of a global economic activity indicator being available in real-time. In emerging and global markets, the Kilian index performs better if extrapolated using the BDI. Nonetheless, in emerging markets, the GEA Tracker, which is based on readily available data from commodity prices, has a better forecasting performance with respect to the Kilian index, even when extrapolated. More important, however, is that we can only find statistically significant predictability using a commodity price-based indicator when there is a careful selection of those commodities. This is true even when performing the selection using LASSO, which shows that, particularly when forecasting financial variables, incorporating information on irrelevant variables deteriorates both the performance as well as the economic gains. The selection of truly relevant variables is therefore crucial. Note, however, that the GEA Tracker, estimated through the genetic algorithm, again outperforms the indicator constructed through LASSO.

Consequently,a mean-variance investor would have significantly profited from using the forecasts provided by the GEA Tracker to construct a portfolio, with a certainty equilavent gain of up to 6,237% in emerging markets. This is proof, not only of the ability of our indicator to forecast the magnitude of the stock returns, but of a better directional accuracy.

5. A Close-Up of the COVID-19 Pandemic

Finally, now that we have carefully examined the properties of the GEA Tracker, we proceed to estimating it a daily frequency for the sample spanned between January 2nd, 2020 to September 15th, 2020.

Daily data for commodity prices is obtained from DataStream and corresponds to the trading price. Not all commodities are traded in stock markets, for which the pool of all commodities is reduced to a total of 37, rather than 47, commodities.[9]

We perform the selection at the end of December 2019, by implementing the genetic algorithm to search for the combination of commodities whose co-movement most closely matches the Kilian index. This is done as in section 2, using monthly frequency data, but with the reduced set of commodities. Real commodity prices are estimated by deflating nominal prices with the U.S. CPI. The selected pool of commodities is kept constant during the entire sample. When the search space is restricted to set Ω_M, the genetic algorithm determines the following group of selected commodities: Henry Hub natural gas, platinum and nickel. In the case of industrial commodities, Ω_I, the genetic algorithm additionally selects the daily price of groundnut oil, and when the search space is extended to include the entire set of commodities, Ω, it rejects the inclusion of any agricultural commodities.

The weight assigned to each commodity price series is determined as in Equation 1. These are estimated monthly, through PCA (for which data are standardized) and are kept constant during the entire month. Daily commodity prices are deflated by the last known value of the U.S. CPI, which is only available at monthly frequency. This lies on the assumption that there are no significant intra-month changes in consumer prices. Finally, daily real commodity prices are standardized using the last estimated monthly mean and standard deviation of the corresponding series.

While the resulting indicators are quite equivalent, for parsimony, we only present the GEA Tracker Metals & Energy, which has proven to have the best properties. This is because it is never subject to extreme idiosyncratic shocks, unrelated to the global economy, of highly

[9]The set Ω_M then includes crude oil, natural gas (US), natural gas (Europe), aluminum, iron ore, copper, lead, tin, nickel, zinc, gold, platinum and silver; Ω_I additionally includes rubber, DAP, urea, coconut oil, groundnut oil, palm oil, soybeans, soybean oil and soybean meal; and Ω also includes barley, maize, sorghum, rice, wheat, orange, beef, chicken, sugar (US), sugar (Europe), sugar (world), cocoa, coffee arabicas and coffee robustas.

volatile commodities.[10]

Figure 8 shows the estimated GEA Tracker.

< Insert Figure 8 about here >

As can be observed in Figure 8, the daily GEA Tracker accurately indicates the evolution of global economic activity during the sample. Particularly, one can observe the large decrease in economic activity occurring in March 11th, 2020, day in which the World Health Organization declared the COVID-19 disease as a global pandemic. This was accompanied by a travel ban set by the United States on Europe and the co-movement of bonds and stocks as investors performed a massive sell-off of their portfolios.

In January and February one can observe a milder slowdown of global economic activity as Chinese premier Li Kegiang urged for decisive efforts to control the pandemic, and several countries implemented travel bans on China. However, in the last week of February, as is reflected by the GEA Tracker, global economic activity came to a stronger halt. On February 21st, the first province in Italy, Lodi, was placed on lockdown and, on February 24th, Asian and European stock indices fell sharply, followed by a plunge of 700 points on the Dow Jones index. This became, at the time, the worst week for stock markets since the Great Financial Crisis.

The first few days of March presented some better news. The seven day decline on global stock indices came to an end, and in March 3rd the U.S. Federal Reserve cut its fed funds rate by 0,5% in an emergency move. The GEA Tracker, therefore, reflects a halt in the decline of global economic activity. However, in March 8th, the Italian primer minister Giuseppe Conte extended the lockdown to the entire region of Lombardy and 14 other northern provinces, as sanitary conditions worsened. March 9th, therefore, was a 'black monday' with the worst drop of the Dow Jones in a single day and the extension of Italy's lockdown to the entire country. This would then be followed by the events of March 11th, after which a number of countries would begin placing drastic mobility restrictions on citizens in an effort to stop the contagion of the coronavirus.

However, as shown by the GEA Tracker, global economic activity slowly began to recover. This was first prompted by fiscal and monetary policy measures in response to the crisis. One can observe the first positive response of the GEA Tracker on March 19th, day in which the U.S. Senate unveiled a $1 trillion economic stimulus package proposal to aid businesses and citizens. This followed the announcement of the Federal Reserve on March 15th of decreasing its benchmark rate for a full 1 percent, and a $700 billion quantitative easing program.

More fiscal and monetary policy actions followed during the month of March, including the British stimulus plan on March 20th, the resumption of construction in China on March 21st, the commitment of the Federal Reserve to buy government-backed securities, and the $2 trillion stimulus package deal from the U.S. Senate on March 25th, among others. News also arrived regarding the lifting of the lockdown in Wuhan.

The recovery of global economic activity continued throughout the rest of the sample, with an observed impulse in May 15th, which is explained by the $3 trillion dollar coronavirus relief

[10]Full results are available upon request.

package that was passed by the U.S. House of Representatives. However, this slow recovery seems to come to a halt in the beginning of August, during which the reactivation of the trade war and tariffs imposed by the U.S. administration further strained trading relationships.

All in all, the GEA Tracker proves to be a reliable and useful real-time indicator of the daily evolution of global economic activity.

6. Conclusions

In this paper, we developed the GEA Tracker, a high-frequency real-time indicator of global economic activity. Not only is this a desirable indicator in normal times for policy-makers and economic agents to perform optimal decisions, but it is paramount during times of economic and social crisis such as the one the world is currently experiencing due to the COVID-19 pandemic.

The GEA Tracker was estimated through a factor model of commodity prices, where commodities were recursively selected through a genetic algorithm to only consider those that are truly relevant to the global economy, at any point in time. We have shown that the issue of the selection of variables is significantly more important than the issue of the dynamics in a factor model. We, therefore, contribute to the strand of the literature that claims that "more is not always better", in which we additionally propose the implementation of the genetic algorithm for the selection of variables. This methodology proves to outperform other methods for selection and shrinkage. Also, in this regard, this is the first paper to empirically examine the recursive selection of variables for the estimation of a factor model.

Moreover, the GEA Tracker proved to have desirable forecasting properties. In fact, when forecasting fluctuations in the Global PMI, we were able to significantly reduce the out-of-sample mean squared error of a benchmark autoregressive model, particularly for the manufacturing sector. We also found significant predictability of stock returns in emerging markets, for which a mean-variance investor would have inexorably profited from using the forecasts provided by the GEA Tracker to weight a portfolio.

Finally, we proved that the GEA Tracker can be estimated at a daily frequency and in real-time, which allows for a closer monitoring of the fluctuations in global economic activity. This becomes critical in times of heightened uncertainty and severe crises.

References

Alquist, R., S. Bhattarai, and O. Coibion (2019). Commodity-price comovement and global economic activity. *Journal of Monetary Economics In press.*

Ang, A. and G. Bekaert (2007). Stock return predictability: Is it there? *Review of Financial Studies 20*(3), 651–707.

Bai, J. and S. Ng (2008). Forecasting economic time series using targeted predictors. *Journal of Econometrics 146*(2), 304 – 317.

Baumeister, C. and L. Kilian (2014). What central bankers need to know about forecasting oil prices. *International Economic Review 55*(3), 869–889.

Bekaert, G. and C. Harvey (1995). Time-varying world market integration. *Journal of Finance 50*(2), 403–444.

Bekaert, G., C. R. Harvey, C. T. Lundblad, and S. Siegel (2011). What segments equity markets? *Review of Financial Studies 24*(12), 3841–3890.

Boivin, J. and S. Ng (2006). Are more data always better for factor analysis? *Journal of Econometrics 132*(1), 169–197.

Campbell, J. Y. and S. B. Thompson (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies 21*(4), 1509–1531.

Clark, T. E. and K. D. West (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics 138*(1), 291–311.

Cooper, I. and R. Priestley (2013). The world business cycle and expected returns. *Review of Finance 17*(3), 1029–1064.

Delle Chiaie, S., L. Ferrara, and D. Giannone (2017). Common factors of commodity prices. *ECB Working Paper 2112.*

Doz, C., D. Giannone, and L. Reichlin (2011). A two-step estimator for large approximate dynamic factor models based on kalman filtering. *Journal of Econometrics 164*(1), 188–205.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*(4), 987–1007.

Herrera, A. M. and S. K. Rangaraju (2019). The effect of oil supply shocks on u.s. economic activity: What have we learned? *Journal of Applied Econometrics 28*(7), 296–324.

Holland, J. H. (1975). *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence.* University of Michigan Press.

Imbs, J. (2006). The real effects of financial integration. *Journal of International Economics 68*(2), 296–324.

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), 1053–1069.

Kilian, L. and X. Zhou (2018). Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance 88*(C), 54–78.

Manescu, C. B. and I. Van Robays (2016). Forecasting the brent oil price: Addressing time-variation in forecast performance. *CESifo Working Paper 6242.*

Markowitz, H. (1952). Portfolio selection. *Journal of Finance 7*(1), 77–91.

Rapach, D., J. Strauss, and G. Zhou (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *Review of Financial Studies 12*(2), 821–862.

Ravazzolo, F. and J. Vespignani (2017). World steel production: A new monthly indicator of global real economic activity. *CAMA Working Paper 42*.

UNCTAD (1990). Review of maritime transport. *United Nations Conference con Trade and Development*.

UNCTAD (2006). Review of maritime transport. *United Nations Conference con Trade and Development*.

UNCTAD (2016). Review of maritime transport. *United Nations Conference con Trade and Development*.

Table 1: R^2_{OS} in percentage (%) values for forecasting the Purchasing Managers' Index (PMI) at the 1-month horizon

	y_t = Global Composite PMI	y_t = Extended Global Composite PMI	y_t = Global Manufacturing PMI	y_t = Extended Global Manufacturing PMI
f_t = Kilian's Index	-1,546	-1,173	-0,275	-1,193
f_t = Kilian's Index (Updated)	-1,967	-2,902	**1,199***	-10,811
f_t = PCA on All Commodities (No Selection)	-0,415	1,179	0,289	**3,708***
f_t = LASSO	1,264	**2,793***	**2,503****	**7,504****
f_t = GEA Tracker, Metals & Energy	2,586	**3,334****	**6,402****	**9,587*****
f_t = GEA Tracker, Industrial	2,934	**3,638****	**6,869****	**10,233****
f_t = GEA Tracker	2,782	**3,879****	**7,282***	**10,850****

* The table above contains the estimated R^2_{OS} percentage values, as described in Equation 12, for the different objective variables, y_t, and global economic activity indicators, f_t, used in the predictive regression model described in Equation 10. *, ** and *** denote significance of 90%, 95% and 99%, respectively.

Table 2: R^2_{OS} in percentage (%) values for stock returns at the 1-month horizon and annualized certainty equivalent gains

y_t	f_t	R^2_{OS}	CE Gain (Unrestricted Weights)	CE Gain (Restricted Weights)
World	Global Composite PMI	-3.136	-9.693	-4.120
	Kilian's (2009) Index	-0,765	-1,340	1,806
	Kilian's (2009) Index (Updated)	1,817	-15,071	1,663
	PCA on All Commodities (No Selection)	-0,566	-1,930	-1,964
	LASSO	0,246	0,723	0,537
	GEA Tracker, Metals & Energy	0,576	-0,544	1,321
	GEA Tracker, Industrial	0,549	-0,771	0,513
	GEA Tracker	0,639	-0,388	0,716
Developed	Global Composite PMI	-1.251	-3.710	-3.018
	Kilian's (2009) Index	0,016	1,187	2,815
	Kilian's (2009) Index (Updated)	-0,841	-4,593	-4,057
	PCA on All Commodities (No Selection)	-2,888	-7,483	-3,212
	LASSO	-0,172	-1,806	-1,588
	GEA Tracker, Metals & Energy	-1,037	-3,812	-3,337
	GEA Tracker, Industrial	-0,565	-2,267	-2,295
	GEA Tracker	-0,538	-1,949	-2,248
Emerging	Global Composite PMI	-2.917	-6.984	-4.644
	Kilian's (2009) Index	-3,119	-5,301	-0,385
	Kilian's (2009) Index (Updated)	**1,261***	-6,361	1,405
	PCA on All Commodities (No Selection)	-0,110	-0,523	-0,260
	LASSO	**1,473****	2,374	4,947
	GEA Tracker, Metals & Energy	**2,364****	5,020	6,237
	GEA Tracker, Industrial	**2,368****	4,967	5,963
	GEA Tracker	**2,268****	4,436	5,272

* The table above contains the estimated R^2_{OS} percentage values, as described in Equation 12, for the different objective variables, y_t, and global economic activity indicators, f_t, used in the predictive regression model described in Equation 14. *, ** and *** denote significance of 90%, 95% and 99%, respectively. Estimations for the certainty equivalent gains, in annualized percentage values, are also reported, as in Equation 19. The estimation is done with no restriction on the portfolio weights, as well as with weights restricted to -0.5 and 1.5 in each period of time.

Figure 1: Classification of commodity prices

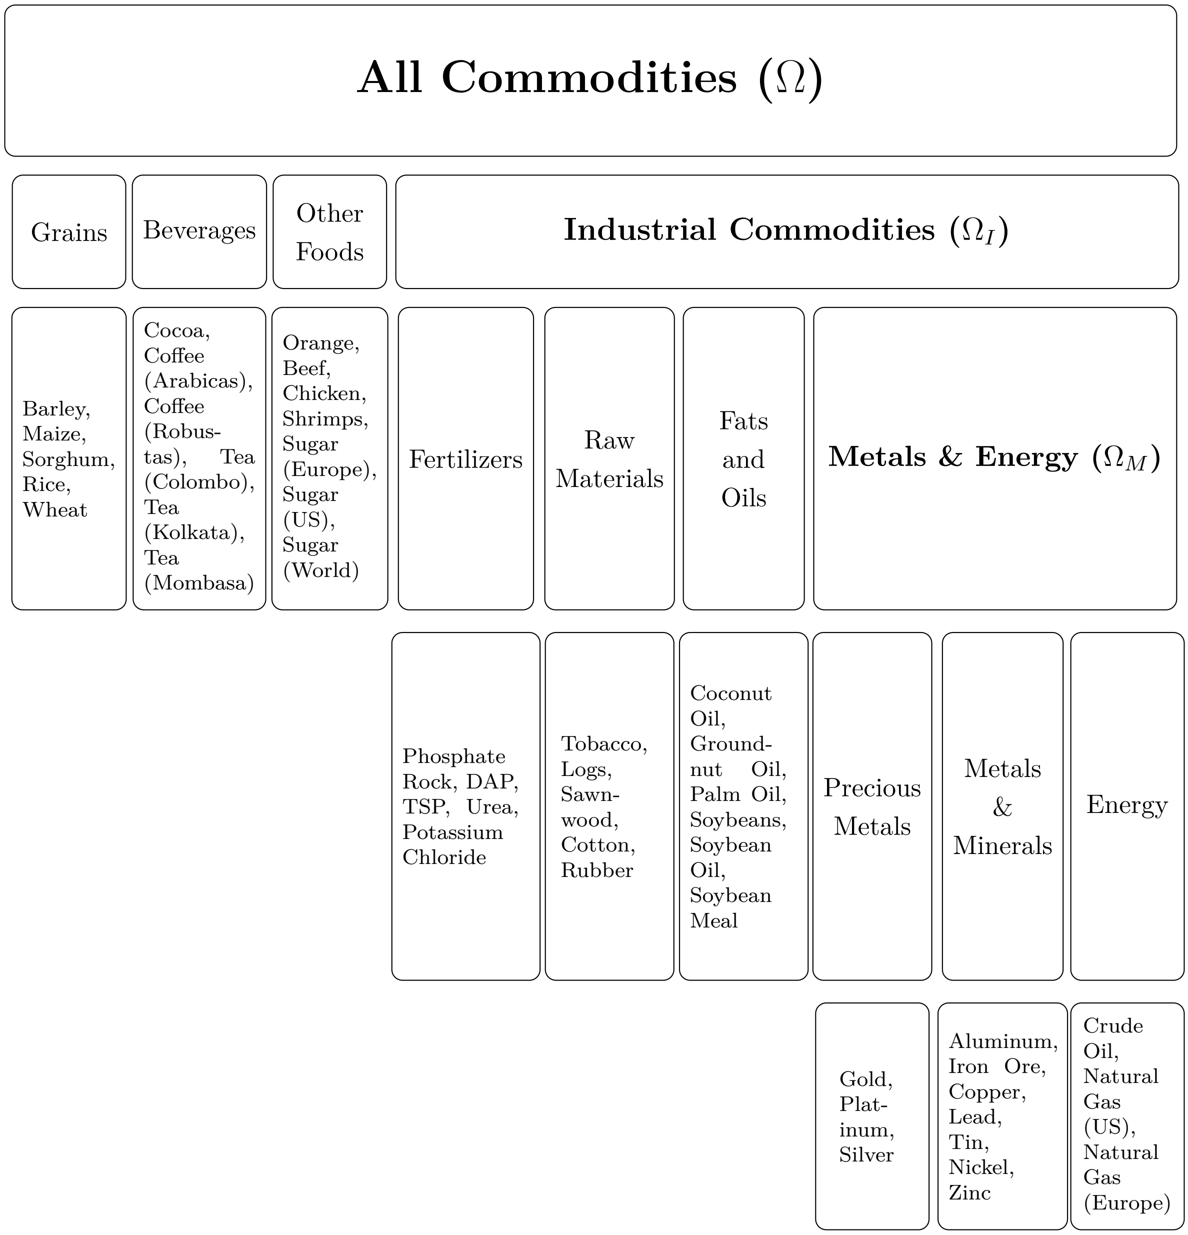

* The figure above shows the classification of commodities used in the estimation of the GEA Tracker. The commodities and group of commodities that pertain to the sets Metals & Energy (Ω_M), Industrial Commodities (Ω_I) and All Commodities (Ω), can be observed.

Figure 2: Recursively estimated indicators of global economic activity

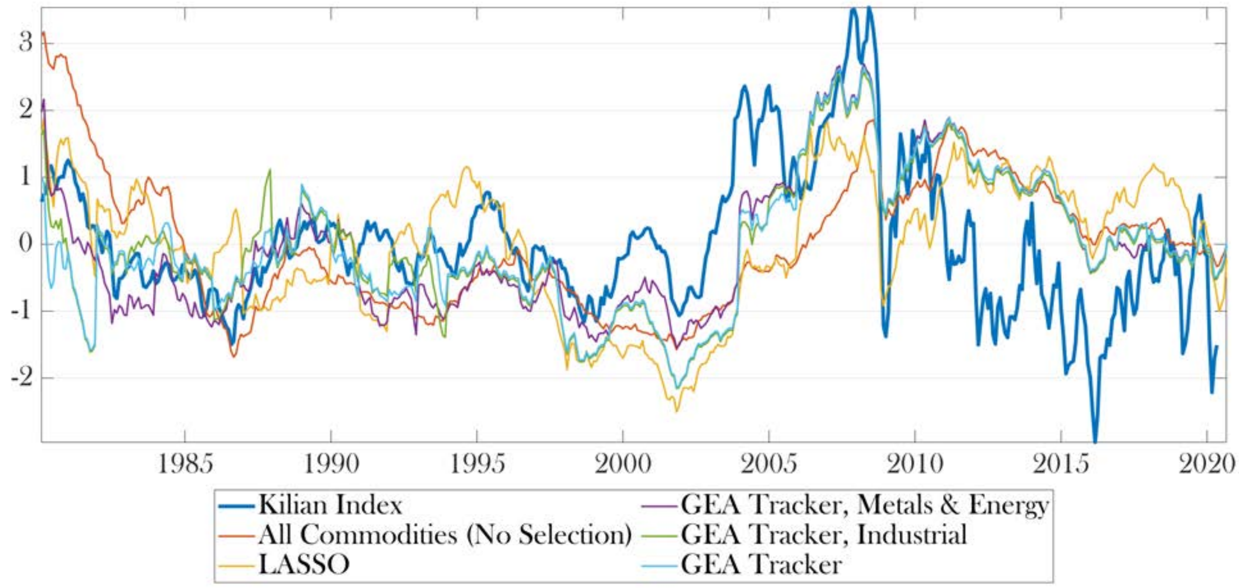

*The figure above shows the benchmark indicator (Kilian, 2009), the indicator generated with no selection, and the indicator estimated recursively by selecting commodities using LASSO, as well as the GEA Tracker Metals & Energy, the GEA Tracker Industrial and the GEA Tracker.

Figure 3: Explained variance of the benchmark indicator by each estimated indicator

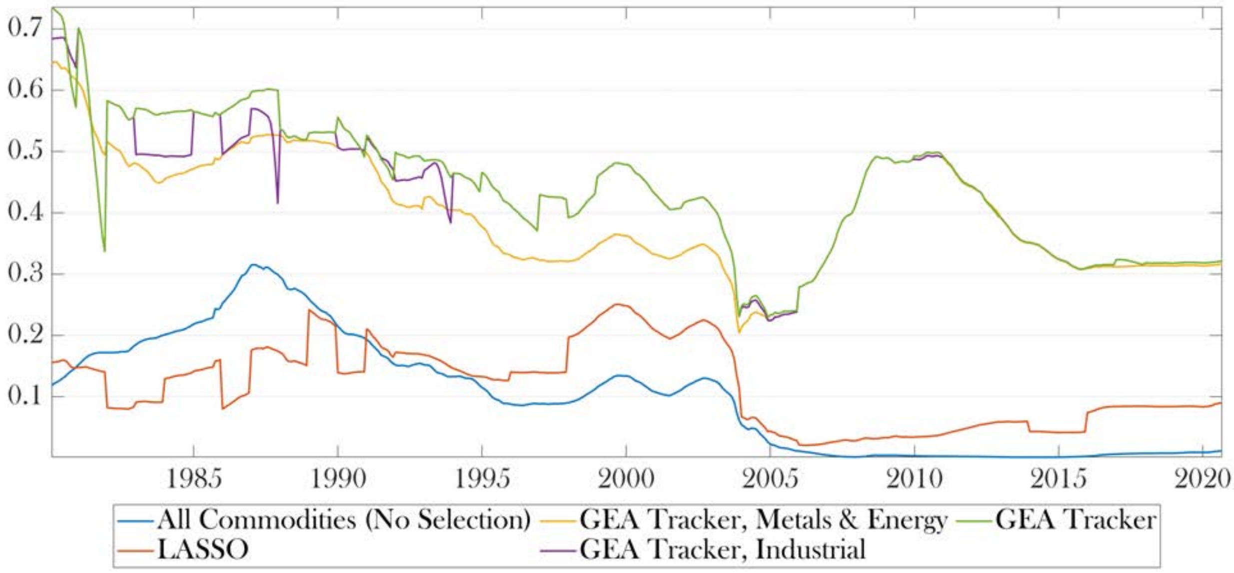

*The figure above shows the variance of the benchmark indicator (Kilian, 2009) explained by each estimated indicator. Each observation corresponds to the R^2 statistic obtained when regressing the benchmark indicator by the first principal component obtained each period t from the historical price series (available until t) of each selected pool of commodities.

Figure 4: Global economic activity indicators

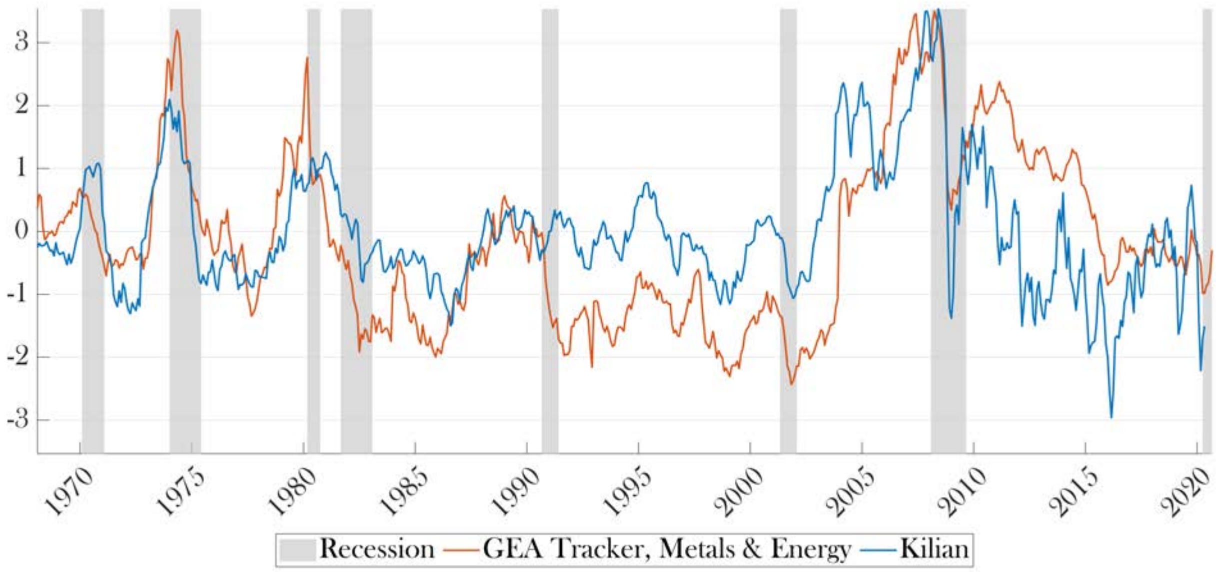

*The figure above shows the GEA Tracker Metals & Energy along with the Kilian (2009) index, for the period 1968M1 to 2020M8. For the GEA Tracker, each observation corresponds to the value estimated with the weights assigned by PCA on the historical log-price series of the commodities selected by the genetic algorithm among Metals & Energy. Light gray shaded areas correspond to U.S. NBER Dated Recessions. The GEA Tracker is estimated in-sample up to 1979M12 and then recursively until 2020M8.

Figure 5: Weights given to each commodity type according to selection method

(a) PCA on all commodities (no selection)

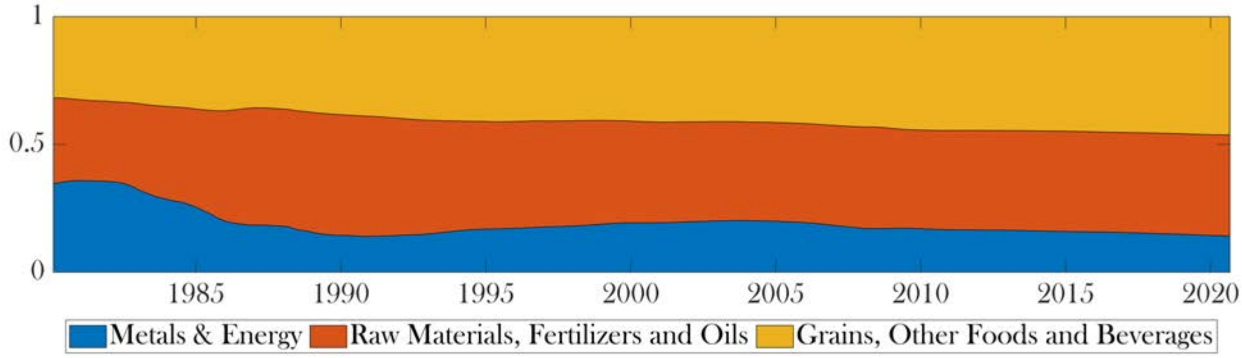

(b) LASSO

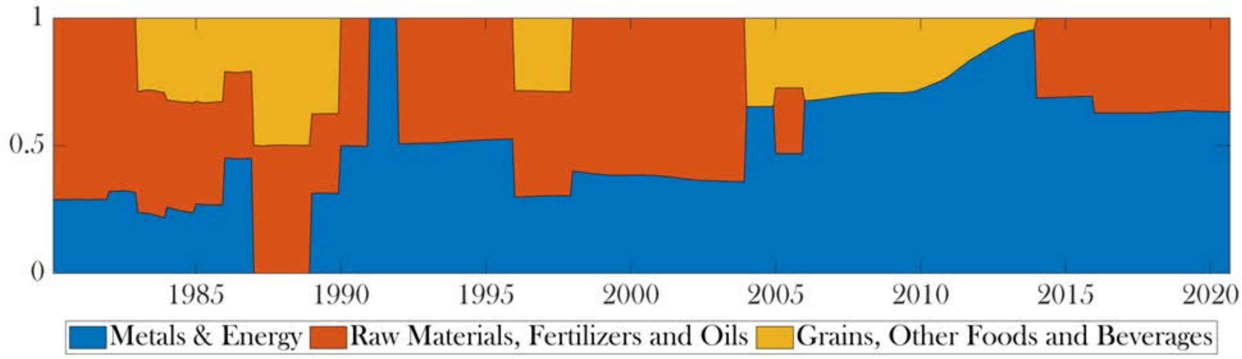

(c) GEA Tracker

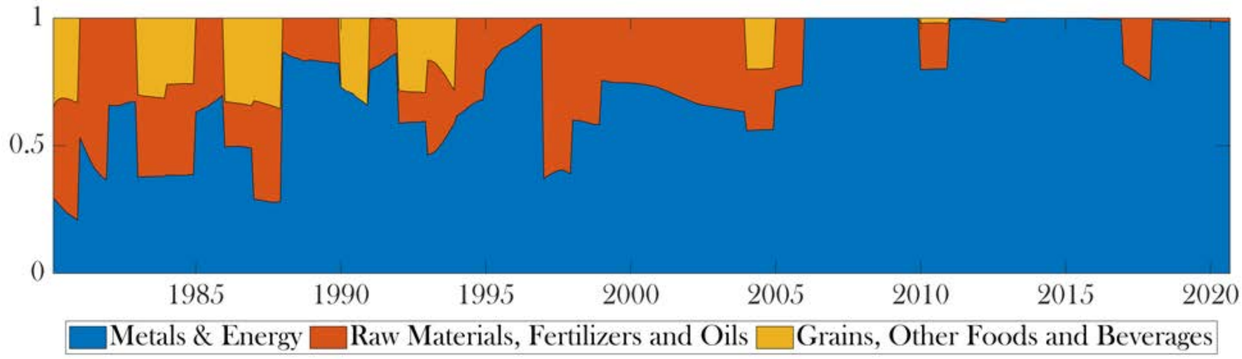

*The figure above shows the time-varying weights given by each estimated indicator to the commodities belonging to the following groups: metals and energy; raw materials, fertilizers and oils; and agricultural commodities. For the case of the indicators estimated with selection (Panels B and C), the weights corresponds to the commodities selected recursively by LASSO and by the genetic algorithm, respectively, with the information available up to date. In Panel A, the pool includes the entire set of commodity prices available from the World Bank.

Figure 6: Weights given to each commodity in the estimation of the GEA Tracker

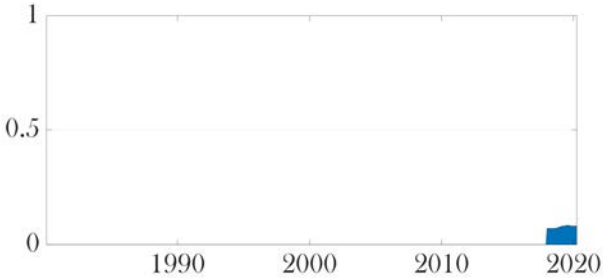

(a) Aluminum

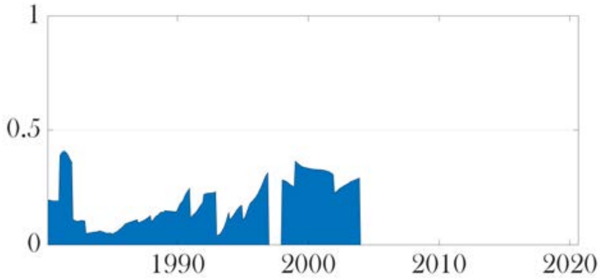

(b) Copper

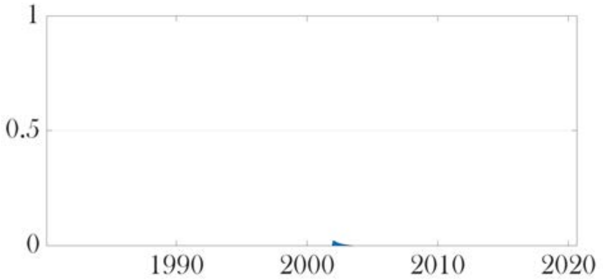

(c) Crude oil

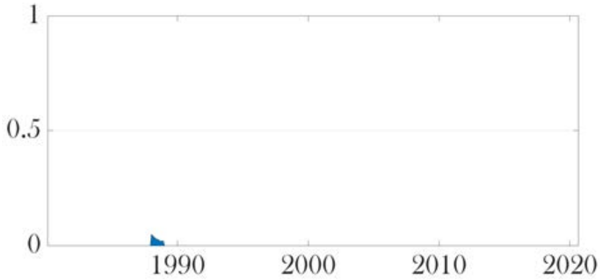

(d) Gold

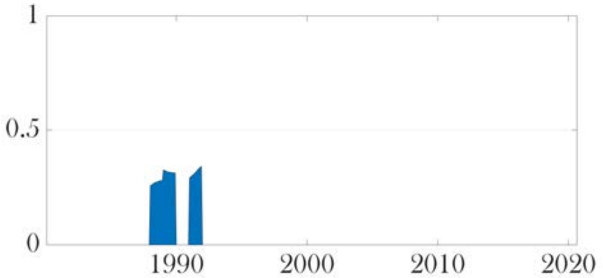

(e) Lead

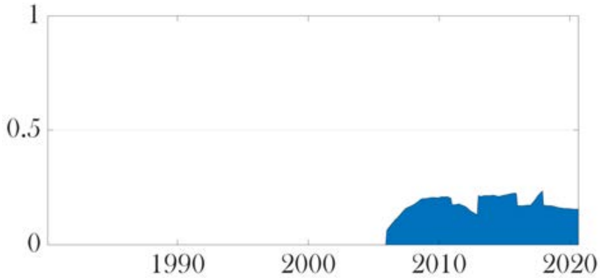

(f) Natural Gas (US)

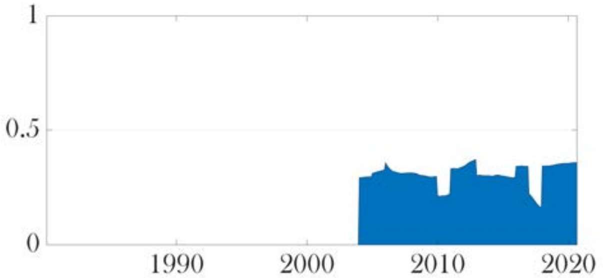

(g) Nickel

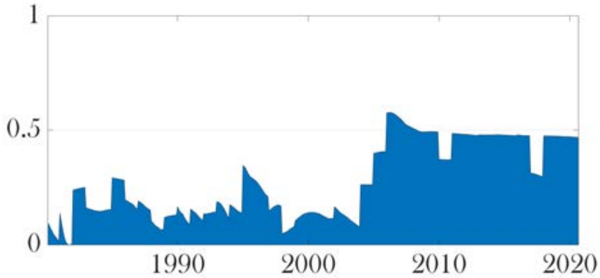

(h) Platinum

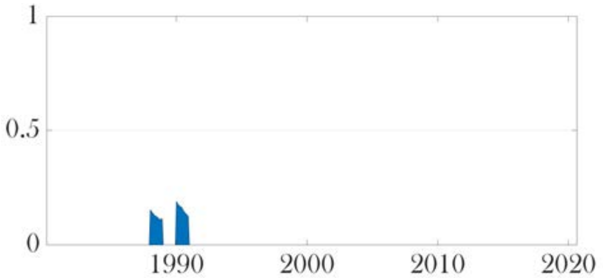

(i) Silver

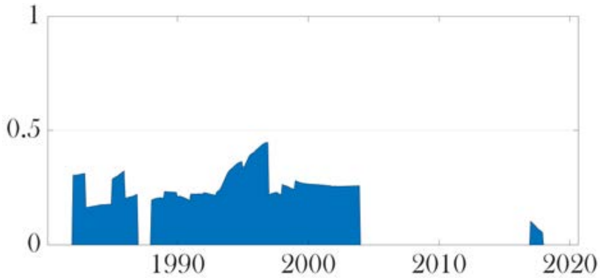

(j) Zinc

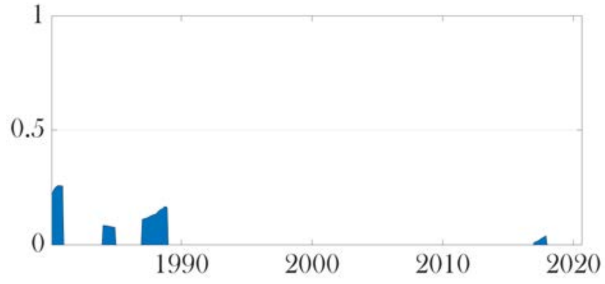

(k) Coconut Oil

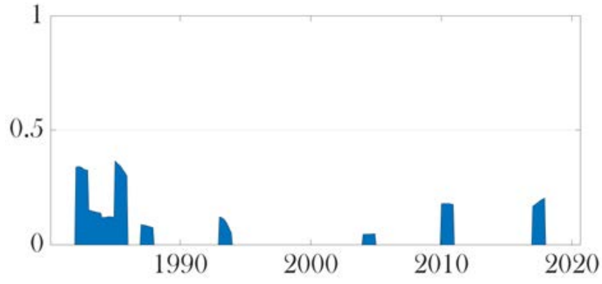

(l) DAP

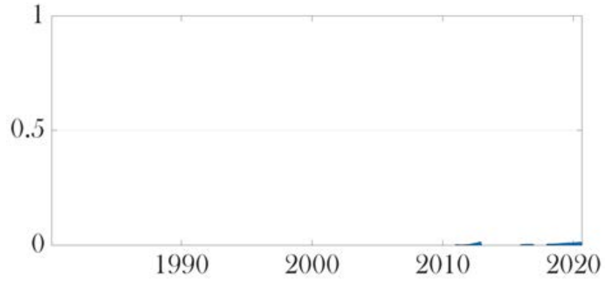

(m) Groundnut Oil

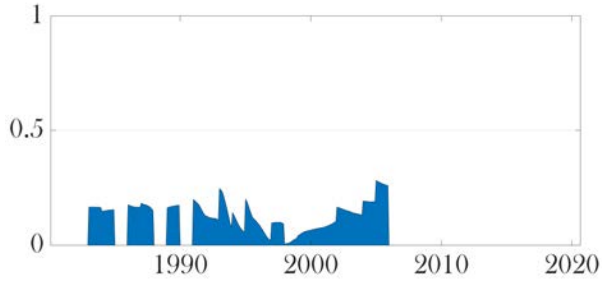

(n) Logs

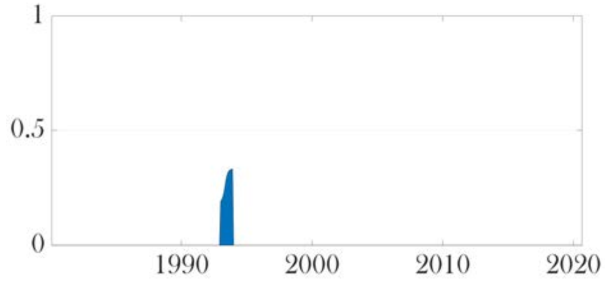

(o) Palm Oil

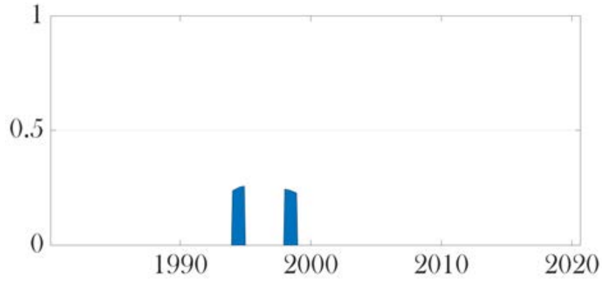

(p) Potassium Chloride

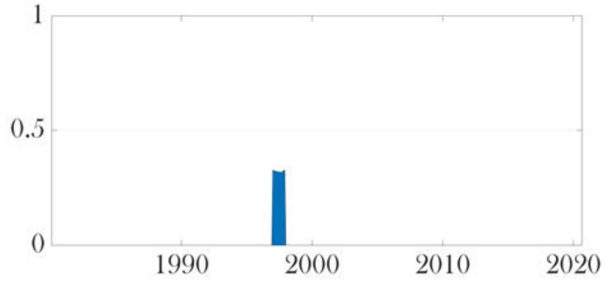

(q) Rubber

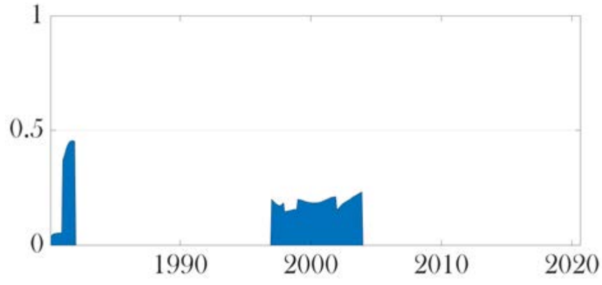

(r) Sawnwood

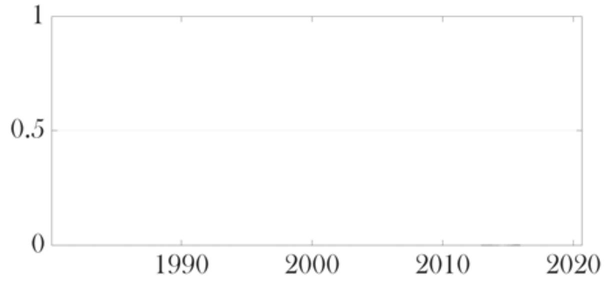

(s) Soybean Meal

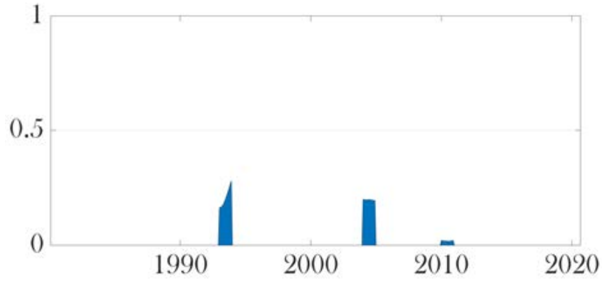

(t) Barley

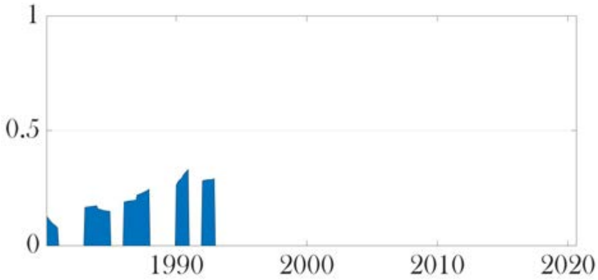

(u) Beef

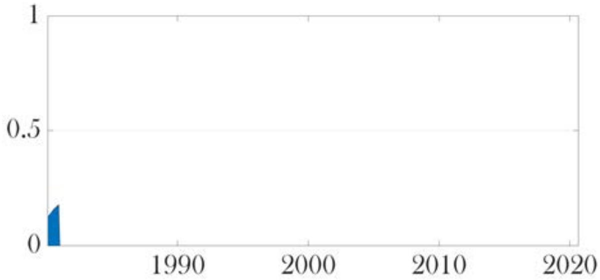

(v) Maize

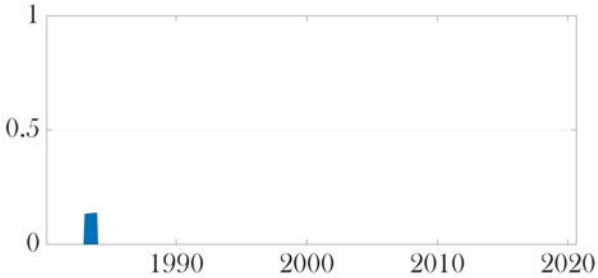

(w) Rice

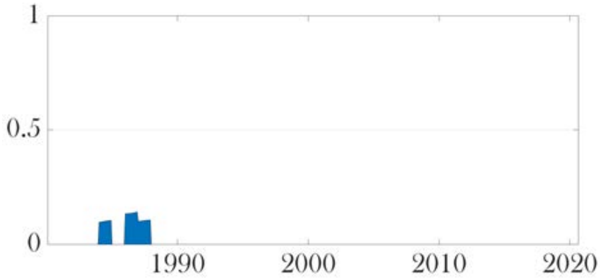

(x) Sugar (US)

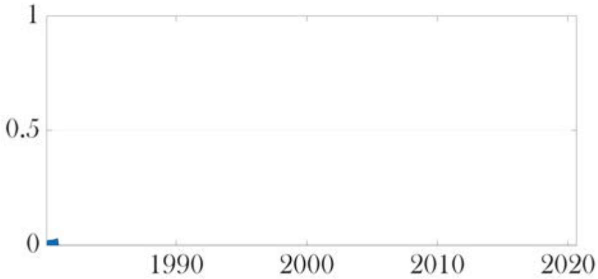

(y) Sugar (World)

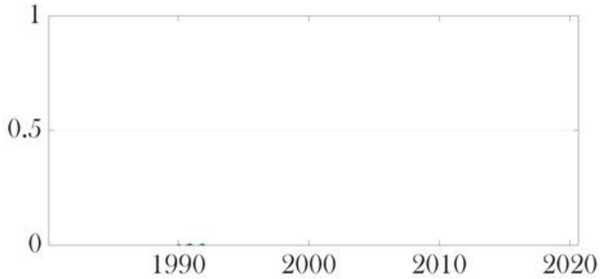

(z) Tea (Mombasa)

*The figure above shows the time-varying weights given by our proposed GEA Tracker to each commodity. We only plot those commodities that are selected by the genetic algorithm at least once during the sample.

Figure 7: Comparison of indicators estimated through dynamic and static factor models, with and without selection

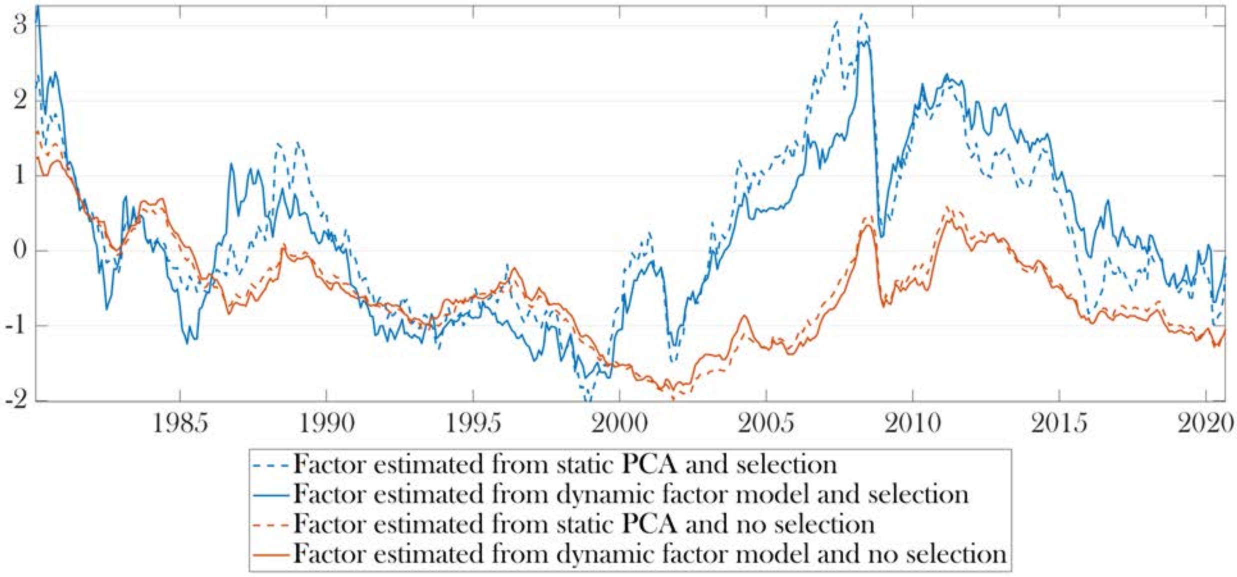

*The figure above shows the in-sample estimation of a global economic activity indicator as the common factor of commodity prices. Depicted in blue, are the indicators estimated by implementing a genetic algorithm to select the commodities that should be included in the estimation. Depicted in orange, are the indicators estimated by implementing no selection at all, but rather including all available commodity price series. Additionally, the indicators estimated through a (static) factor model, as the one described in Equation 1, are represented by dashed lines; whereas the indicators estimated through a dynamic factor model, as described in Equations 1, 8 and 9, are represented by continuous lines.

Figure 8: Evolution of global economic activity in year 2020

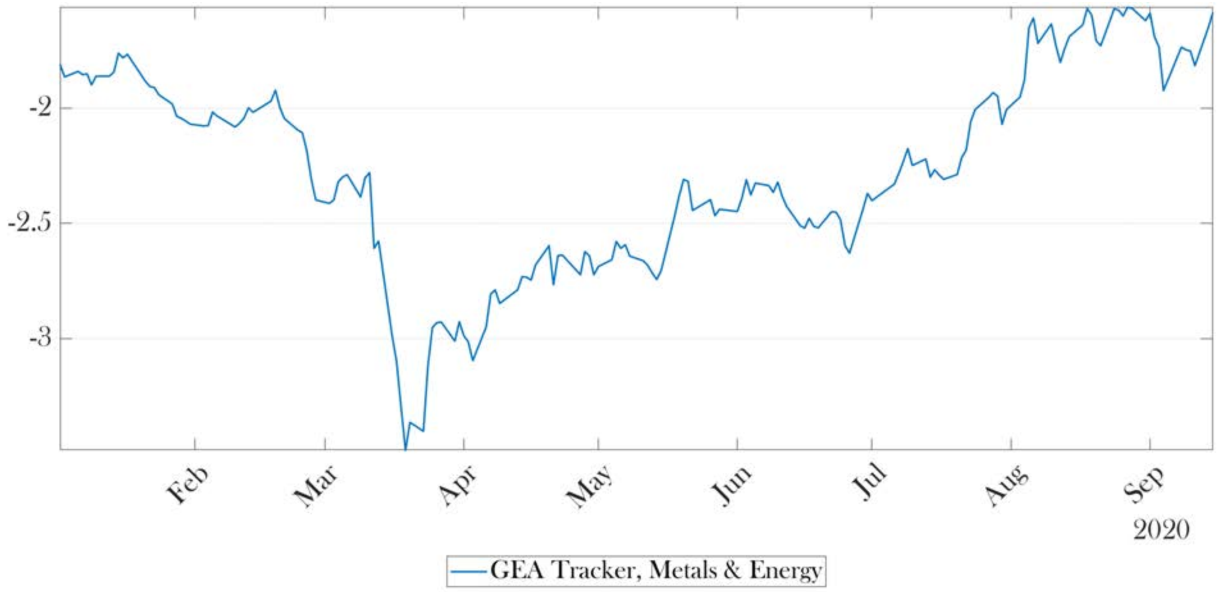

*The figure above shows the GEA Tracker in daily frequency for the sample spanned from January 2nd, 2020 to September 15th, 2020.

Appendices

A. A Genetic Algorithm for the Selection of Commodities

For the selection of commodities to estimate the GEA Tracker, we begin by defining an individual A_j as a genome of n binary genes $a_{ij} \exists 1, 0 \; \forall \; i \exists \{1, \ldots, n\}$.

$$A_j = (a_{1j}, a_{2j}, \ldots, a_{nj})$$

where n is the total number of commodities and a_{ij} is assigned the value of 1 if commodity i is included in the estimation of the GEA Tracker, and 0, if it is not.

We then define a population as a set of J individuals, A_j, $\forall \; j \exists \{1, \ldots, J\}$. A population P is therefore a matrix of size $J \times n$ where each row corresponds to the genome of individual A_j.

The genetic algorithm will evolve the population through a total number of generations or iterations, K. The population P_{k+1} is created from population P_k, where $k \exists \{0, \ldots, K-1\}$. This is done as explained followingly.

A.1. Initial Population

For the initial population P_0, every gene $a_{ij} \; \forall \; i \exists \{1, \ldots, n\}$ and $j \exists \{1, \ldots, J\}$, is randomly generated through a uniform distribution. This is done by assigning 0 or 1 to each element a_{ij}, following a binomial distribution where the probability of $a_{ij} = 1$ is 0.5. Otherwise, $a_{ij} = 0$.

A.2. Fitness Function

Once a population is generated, the fitness values, F, of the population are estimated. F is a $J \times 1$ size vector, in which each element is estimated as follows:

(i) For each A_j, a matrix X_j of real commodity price time series is constructed. Each column of X_j contains the price series of commodity i if $a_{ij} = 1$.

(ii) Matrix X_j is standardized and Principal Component Analysis is estimated to extract the first principal component, denoted f_j.

(iii) f_j is used as a regressor in the following regression model:

$$f_t^* = \mu + \beta f_{jt} + \epsilon_t \qquad \epsilon_t \sim N(0, \sigma^2)$$

where f_t^* is the Kilian (2009) index at time t, μ is a mean constant, β is a slope-coefficient, and ϵ_t is an error term with a normal zero-mean distribution and variance σ^2. The regression is estimated using Ordinary Least Squares (OLS).

(iv) The coefficient of determination R^2 is then determined as the proportion of the variance in the dependent variable that is explained by the regressor.

$$R^2 = 1 - \frac{\sum_t \epsilon_t^2}{\sum_t (f_t^* - \overline{f_t^*})^2}$$

(v) The fitness value for individual A_j is defined as the estimated coefficient of determination R^2.

(vi) Furthermore, for all individuals A_j in a population where the genome contains 2 or less genes, a_{ij}, equal to 1, the fitness values are penalized and set to -99999. This is to avoid the selection of less than 3 variables, in which PCA degenerates.

The estimated fitness values of all individuals A_j of the population then constitute the vector F.

A.3. Scaling Function

The scaling function $g(\cdot)$ then converts the fitness values F into values that are suitable for the following steps of the algorithm. To do so, the population individuals A_j are first ranked from highest to lowest according to their fitness value. The scaled values are given by

$$g(A_j) = \frac{1}{\sqrt{r_j}} \qquad \forall\, j \exists \{1, \ldots, J\}$$

where r_j is the rank of individual A_j. A vector F_S is then defined as $F_S \equiv g(F)$ and contains the scaled fitness values of each individual A_j.

A.4. Creating Population P_{k+1}

Given the population P_k and its scaled fitness values, we generate population P_{k+1} by performing three operations: elitism, crossover and mutation. The elite function allows for the most fit individuals of P_k to be included in P_{k+1}, while crossover and mutation are reproduction functions that combine the genetic information of individuals in P_k to generate new individuals that will be part of P_{k+1}. These new individuals are referred to as children, whereas the individuals from P_k from which they are created are called parents.

P_{k+1} will then be constructed as the union of the set of elite individuals denoted R_e and the set of "crossovered" and "mutated" children, denoted R_c and R_m, respectively.

1.4.a. Elite Function

An elite number of individuals with the best fitness values are guaranteed to survive to the next generation P_{k+1}. The elite count J_e is set to be 5% of the population size, J.

$$J_e = 0.05J$$

The elite individuals are those who have the J_e highest fitness values and create the elite set R_e. R_e is then a matrix of size $J_e \times n$ that contains the genomes of the most fit individuals.

1.4.b. *Reproduction Functions*

80% of the children will be created through crossover, while the remaining 20% will be created through mutation. Therefore, to keep a constant population size throughout all generations $k \exists \{0, \ldots, K\}$, the number of children to generate is estimated as follows:

$$J_c = round(0.80(J - J_e))$$
$$J_m = J - J_e - J_c$$

where J_c is the number of children generated through the crossover function, and J_m is the number of children generated through the mutation function.

Therefore, the total amount of parents, J_p, from population P_k we need is estimated in the following way:

$$J_p = 2J_c + J_m$$

Then, for reproduction, we first implement a selection function on P_k in a way that the most fit individuals have a higher probability of being parents.

1.4.c. *Selection Function*

The selection of parents from the current population P_k is performed as a stochastic uniform selection. This operation consists on drawing a line in which each section of the line corresponds to an individual A_j of the population. Every section has a length proportional to its scaled fitness value. We then move along the line in steps of equal size, selecting the individual it lands on as a parent. The size of the steps is set to $\frac{1}{J_p}$ and the first step is given by a randomly generated number in the range $[0, \frac{1}{J_p}]$.

The most fit individuals will have a higher probability of being selected as parents, with their genetic information being passed on to the following population P_{k+1}. Note that any given individual of the population can be chosen to be a parent more than once.

The set of parents is then a matrix of size $J_p \times n$ which contains the genome of each selected individual.

1.4.d. *Crossover Function*

The crossover function we use is the scattered crossover. This function generates a random binary vector of size $1 \times n$, with each element corresponding to gene a_i. This random vector indicates whether the child will inherit the value for each gene a_i from parent A (for values equal to one) or from parent B (for values equal to zero). We illustrate this using a simplified example. Suppose parents A and B have the following genetic information:

$$Parent\ A = [A\ B\ C\ D\ E]$$
$$Parent\ B = [a\ b\ c\ d\ e]$$

If the random binary vector were, [1 1 0 1 0], then the child would be defined in the following way:

$$Child = [A\ B\ c\ D\ e]$$

The crossover function is applied J_c times, to a total of $2J_c$ parents from the set of selected individuals to create a set of R_c "crossovered" children. The pair of parents for each child is selected randomly. R_C is then a matrix of size $J_c \times n$ which contains the genome of all "crossovered" children.

1.4.e. Mutation Function

The gaussian mutation function adds a random number taken from the Gaussian distribution $N(0, \sigma_k)$ to each gene a_{ij} of the parent vector, where σ_k is determined by

$$\sigma_k = \sigma_{k-1}(1 - \tfrac{1}{J})$$

Where σ_k is shrunk at each generation k, reducing the probability of mutation as the algorithm approximates the optimal solution. Each a_{ij} is then set to its closest value in the set $\{0, 1\}$.

The mutation function is applied to a total of J_m parents to create a set R_m.

R_m is then a matrix of size $J_m \times n$ which contains the genome of all "mutated" children.

A.5. Replacement

Finally, the new generation P_{k+1} is created by the union of the elite set of children, R_e, the crossovered children, R_c, and the mutated children, R_m. Fitness evaluation, selection, elite generation, reproduction and replacement are then repeated K times on all populations P_k, where $k \exists \{1, \ldots, K\}$, until a stop selection criterium is met.

A.6. Stop Selection Criteria

The algorithm repeats until either:

(i) The maximum number of generations, K, is achieved, which is set to be $100\times$ number of selection variables n.

(ii) The function tolerance, defined as the average relative change (from k to $k+1$) in the fitness function value, reaches a maximum objective value. In our application, this is specified as 1×10^{-6}.

Acknowledgements

Special thanks to Fernando Perez de Gracia, Juan Carlos Molero, Francesco Ravazzolo, Christiane Baumeister, and Junsoo Lee for many insightful comments and feedback on earlier versions of the paper. We also thank useful comments received during the 38th International Symposium on Forecasting (2018), the IV Navarra-Basque Country Macroeconomics Workshop (2017), and during seminars at the Free University of Bozen, ICADE (Universidad Pontificia Comillas) and Universidad de Navarra. We are also grateful to Romain Aumond for his extremely valuable research assistance. Finally, Elena Diaz recognizes financial aid from the “Programa de Ayudas de la Asociacion de Amigos de la Universidad de Navarra” and Banco de Santander. The views expressed here are those of the authors and do not represent the views of the ECB or the Eurosystem.

Elena Maria Diaz
Universidad Pontificia Comillas, Madrid, Spain; Universidad de Navarra; email: emdaguiluz@comillas.edu

Gabriel Perez Quiros
European Central Bank, Frankfurt am Main, Germany; CEPR; email: gabriel.perez_quiros@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB’s website.

PDF ISBN 978-92-899-4451-9 ISSN 1725-2806 doi:10.2866/961578 QB-AR-20-157-EN-N