QUMS WORKING PAPER SERIES
(Finance)

http://www.qub.ac.uk/schools/QueensUniversityManagementSchool/Research/ResearchandResearchers/WorkingPapers/ByGroup/

The Baltic Dry Index: Cyclicalities, Forecasting and Hedging Strategies

Fotis Papailias and Dimitrios D. Thomakos

Queen's University Belfast and University of Peloponnese

Working Paper FIN 14-12

QUEEN'S UNIVERSITY MANAGEMENT SCHOOL
WORKING PAPER SERIES
Queen's University Belfast
185 Stranmillis Road
Belfast BT9 5EE

October 2014

The Baltic Dry Index: Cyclicalities, Forecasting and Hedging Strategies

Fotis Papailias* Dimitrios D. Thomakos†

Abstract

The cyclical properties of the annual growth of the Baltic Dry Index (BDI) and their implications for short-to-medium term forecasting performance are investigated. The BDI has a cyclical pattern which has been stable except for a period after the 2007 crisis. This pattern has implications for improved forecasting and strategic management on the future path of the BDI. To illustrate the practicality of our results, we perform an investment exercise that depends on the predicted signs. The empirical evidence supports the presence of the cyclical component and the ability of using forecast signs for improved risk management in the freight sector.

Keywords: Baltic Dry Index, Commodities, Concordance, Cyclical Analysis, Forecasting, Freights, Hedging, Turning Points.

***Corresponding Author.** quantf research, www.quantf.com. Queen's University Management School, Riddel Hall, 185 Stranmillis, BT9 5EE, Belfast, Northern Ireland.`f.papailias@qub.ac.uk`, `f.papailias@quantf.com`; Tel.: +44-28-9097-4667; Fax: +44-28-9097-4201

†quantf research, www.quantf.com. Department of Economics, University of Peloponnese, Greece and Rimini Centre for Economic Analysis, Rimini, Italy. `thomakos@uop.gr`, `thomakos@quantf.com`

1 Introduction

The dry bulk shipping sector has long been of interest to investment banks, institutions and academia. The size-specific indices like Supramax, Panamax and Capesize along with the overall Baltic Dry Index (BDI) are being observed daily by economists and market investors. The BDI is defined as a daily weighted average freight price to ship raw materials across the globe in order to be used in the production process. Therefore, it incorporates aspects of the future economic activity and thus has the characteristics of a leading economic indicator. This economic linkage may provide a misleading idea of order regarding the market dynamics of freight rates, as there are at least two typical problems in this market that complicate its dynamics: endogeneity and supply lags. It is straightforward to see that shipping costs affect and are affected by global activity, which means that is difficult to reliably use even current activity estimates to foretell the path of freight rates. Regarding the supply lags, the supply of freight services (i.e. new and second-hand ships, scrapping etc.) is very inelastic (reacts poorly/slowly to price changes) as it is limited by supply lags and often lacks of market depth, while the demand for freight services tends to be very elastic. Generally, such a set up, favours the eventual resolution of supply/demand imbalances and smooths out the path of the BDI.

The transportation of dry bulk goods affects a variety of markets and not just the shipping market. Items such as coal, steel, iron ore, foodstuffs such as corn, wheat and many others, indicate that BDI variation should have strong association with the commodities market as well. Recent professional research indicates that this dual causality of BDI and commodities plays a key role in the pricing process. For example, Bornozis (2006) sheds some extra light regarding the global factors that affect the supply and demand in this sector while a report by Giannakoulis and Slorer (2012) reports that the daily crude steel run rate for February, 2012 and iron ore imports were surprisingly high. A report by Nomura Equity Research (2012) further mentions excess supply issues with 2012 being the third consecutive year of double-digit supply growth while demand has never recorded a double-digit growth historically. It is important to notice that these reports expect the BDI to rebound from its current levels.

The academic literature on the BDI and various sub-indices, and the shipping freight rates in general, has a long history and many papers have analysed various aspects of the behaviour and time series properties of these indices. Driehuis (1970) is among the first to provide a very thorough investigation of the liner freight rates, including a well-formulated economic-theoretic model. Marlow and Gardner (1980), also have an early model on the dry bulk shipping sector and Beenstock and Vergottis (1989a) and Beenstock and Vergottis (1989b) build an econometric model for the world tanker market and the dry bulk market. In contrast to Adland and Cullinane (2006), Koekebakker, Adland and Sodal (2006) and Batchelor, Alizadeh and Visvikis (2007) are interested in the BDI series as a whole rather than analysing the spot or forward rates separately. More recently, Goulas and Skiadopoulos (2012) analyse the efficiency of the IMAREX futures markets and Lin and Sim (2013) investigate the relationship and effect of the BDI with trade and income. Comparisons of volatility in the dry-cargo ship sector have been conducted by Kavussanos (1996), while the seasonal properties and forecasting in the dry bulk shipping sector are investigated in Cullinane, Mason and Cape (1999), Kavussanos and Alizadeh (2001) and Kavussanos and Alizadeh (2002).

Using the work of the last mentioned authors as our point of departure, in this paper we are mainly concerned with the cyclical (and not just the seasonal) characteristics of the BDI annual growth series. We test for the cyclical properties of the series, with special emphasis on the frequency part that is closer to the business cycle, and we develop different models to capture and interpret this characteristic. Our analysis shows that there is a strong cyclical pattern of cycle duration of between 3 to 5 years and that this pattern holds quite well across different sub-samples. The spectrum of the series shows distinct peaks at the corresponding frequencies and even a simple trigonometric regression using these frequencies shows a relatively good fit. Furthermore, the cyclical pattern we find affects the way the mean and persistence operate in the data: using a functional model we can actually provide a novel aspect at looking at the way that BDI growth moves along the cycle path. After our cyclical analysis we consider the problem of forecasting the BDI annual growth series.

Here we take a quite comprehensive stance, compared to the existing BDI-related literature, and we consider a variety of models and model averages that incorporate explanatory variables – carefully selected by top-down elimination – and the cyclical component. Our forecasting exercise is conducted with a focus on the short-to-medium term horizon and we evaluate one, six and twelve months-ahead forecasts of the BDI growth. Our results indicate that a considerable proportion of BDI growth variation can be predicted by a combination of explanatory factors and the cyclical pattern that exists in the series. Finally, our forecasting experiments and evaluation further improves the work of Denning, Riley and DeLooze (1994).

If the above results are considered as a point for further analysis, a reasonable question to ask next is whether we can use our models for maritime risk management purposes. That is, if cyclicality is indeed present in the data and a model can capture it then there might be a way of using this, for example, in hedging the path of the BDI or having a portfolio of other assets replicating it or even speculating on its future performance. Investment banks, shipping firms and individual investors that make their business choices based on expectations about the BDI could benefit in terms of correct model timing. We thus go a step further and show how our forecasting results can be put in real-life use in evaluating a straightforward investment strategy: this strategy compares the performance of model-based investment decisions against some alternative benchmarks. The results from this approach indicate that the timing ability of our suggested models works well in an investment-decision context and can thus be further exploited for risk management purposes. The same could be used for freight futures trading.

The rest of the paper is organised as follows: in Section 2 we discuss our data and their descriptive statistics; in Section 3 we present the tests, analysis and discussion on the cyclical behaviour of the BDI annual growth; in Section 4 we first introduce the different models we use for predicting the BDI growth, we evaluate and discuss their forecasting performance and we present the results of the BDI investment strategy that uses the previously described forecasting models; in Section 5 we offer some concluding remarks.

2 Data

In our analysis we use monthly data for the BDI and a number of related variables. The full sample range, after adjusting for the computation of annual growth rates, spans from January, 1993 to March, 2012 for a total of $T = 231$ monthly observations. The other variables that we consider include: commodities such as CORN, COTTON and WHEAT, COPPER, CRUDE AND TIN; the Morgan Stanley global indices for emerging and developed markets MSCIEM and MSCIDEV; the British pound/U.S. dollar exchange rate, GBPUSD, and the dollar index DXY; and, finally, the SPREAD, which denotes the difference between the 10-year and 2-year US Treasury yields. The whole dataset is collected from Bloomberg and is expressed in annual growth as well. The basic descriptive statistics of the annual growth series are given in Table 1. In the first panel of the table we present the statistics for each series and in the second panel of the table we have the contemporaneous, full sample, correlations of BDI with the other variables.

The statistics present some interesting features. The BDI has the highest average growth among all other variables, except the SPREAD, and this is due to the large increase that it exhibited before the financial crisis in 2008 – and also due to the large decrease after the crisis. Of comparable, although smaller, magnitude is the average growth of CRUDE, COPPER, TIN, MSCIEM and CORN – note that the corresponding standard deviations are half that of BDI for all these variables. Turning to the correlations, we can see that – in absolute magnitude – the highest correlations are (positive) for the MSCIEM, GBPUSD and (negative) for the DXY index, at about 50%. Note that they all make sense, in that a weaker US dollar was associated with the period of higher global growth, thus higher MSCIEM growth, and the large increase in the BDI. After these three variables we see that COPPER and TIN follow with (positive) correlations of 46% to 48%, and then we have a group of variables with positive correlations of close to 30%. Of interest is also the low correlation that is exhibited between BDI and CORN and BDI and SPREAD – as we will see, however, in the forecasting exercise we conduct later on the low contemporaneous correlation of BDI and CORN does not prevent CORN as being the single explanatory variable that performs best in predicting BDI variation.

For illustration purposes, consider Figure 1 and Figure 2. There we compare the BDI to all other variables of interest in actual levels and in annual percentage change (annual return). We see that, although some of the variables seem to deviate from the BDI, they tend to move in the same direction before 2007 and after 2010. In the cyclical analysis that follows we try to formalise what can be casually seen from these two figures.

3 Cyclical Analysis of the BDI Annual growth

3.1 Identification of Turning Points & Tests of Synchronicity

In our analysis we consider the results in Harding and Pagan (2006), where a coherent methodology is presented for testing cycle synchronicity. The testing methodology proposed therein presupposes that one has available indicator variables that identify expansion and contraction periods for each series. There are various ways of getting these indicator variables but here we follow a straightforward approach as presented in Harding (2008). We briefly summarise the methodology below while full details can be found in the above papers.

Consider a time series of interest y_t and suppose that we would like to find its local turning points (local maxima and minima) in a window of k observations. Then, these local peaks and troughs are given by,

$$\begin{aligned} \wedge_t &\stackrel{\text{def}}{=} I\left[(y_{t-k}, \ldots, y_{t-1}) < y_t > (y_{t+1}, \ldots, y_{t+k})\right] \\ \vee_t &\stackrel{\text{def}}{=} I\left[(y_{t-k}, \ldots, y_{t-1}) > y_t < (y_{t+1}, \ldots, y_{t+k})\right] \end{aligned} \tag{1}$$

where $I(\cdot)$ is the indicator function. While these two variables can be used to mark expansions and contractions they have the problem that cycle phases may not alternate and, to alleviate this problem, a form of censoring can be used. To do so one uses the following recursion to construct a single binary variable that marks expansions and contractions and has the cycle phases alternating,

$$S_t \stackrel{\text{def}}{=} S_{t-1}(1 - \wedge_{t-1}) + (1 - S_{t-1}) \vee_{t-1} . \tag{2}$$

Based on the above series the alternating turning points are then given by,

$$\begin{aligned} \wedge_t^a &\stackrel{\text{def}}{=} S_t(1 - S_{t+1}) \\ \vee_t^a &\stackrel{\text{def}}{=} (1 - S_t)S_{t+1} \end{aligned} \tag{3}$$

The focus of the analysis is then in the S_t series. Consider two such series S_{tx} and S_{ty} for two underlying variables X_t and Y_t; where Y_t denotes the BDI annual percentage change series and X_t denotes another variable which may be commodities, foreign exchange rates and so on. Let $\rho_S \stackrel{\text{def}}{=} \mathsf{Corr}\left[S_{ty}, S_{tx}\right]$ denote the correlation coefficient between the S_{tx} and S_{ty} series. Following Harding and Pagan (2006), the series are said to be in strong positive synchronisation when the following conditions hold,

$$SPS : \mathsf{E}\left[S_{ty} - S_{tx}\right] = 0 \text{ and } \rho_S \neq 0 \tag{4}$$

where if in addition $\rho_S = 1$ then we have the series to be in strong perfect positive synchronisation. On the other hand, we have that the series are in strong negative synchronisation if they have zero correlation, i.e. when we have,

$$SNS : \rho_S = 0 \tag{5}$$

without the need to consider the properties of the mean difference $\mathsf{E}\left[S_{ty} - S_{tx}\right]$. Testing the above conditions is easily done via a GMM approach that accounts for the presence of potential heteroscedasticity and autocorrelation. Our results refer to testing these two hypotheses and are all summarised in Table 2. Note that, in addition to the estimates and their z-statistics, we report the (estimate of the) concordance index C which relates to the correlation ρ_S and measures the proportion of time that the two series are in the same phase. Additional details about the structure of the tests and the concordance index can be found in Harding and Pagan (2006). We use two values for k, one corresponding to an annual cycle ($k = 6$ months on either side of the turning point) and one corresponding to a 5–years cycle ($k = 30$ months).

3.2 The Turning Points and Spectrum of the BDI Annual Growth

We first look at the visual characteristics of the BDI series and its growth. The peak and trough points are estimated using equation (3). The cyclical features are evident in Figure 3, both for the level of the series and its annual growth. Notice that for almost a decade (1993 to 2001) there was an almost deterministic cyclical pattern since peaks and troughs occur in similar values for both series and are about equally spaced (notice that this was about the period that the two papers of Kavussanos and Alizadeh (2001) and Kavussanos and Alizadeh (2002) have used in their analysis). From 2001 onwards it appears that the duration of the cycle pattern has increased and there is a break in the systematic seasonal behaviour, although as we see later the longer-term cyclical behaviour is still there.

Prying a bit more into the behaviour of the peaks and troughs on the annual growth series, we estimate that, for $k = 6$, the average amplitude during the expansion part of the cycle was about 0.75% (percentage points) while the average amplitude during the contraction part was about -0.85%. On the other hand, for the larger – and more relevant cycle – of $k = 30$ there is considerable asymmetry in these amplitudes as they were estimated to be 1.18% and -3.19% respectively. These numbers are for the full sample, that includes the large fluctuations of 2002 onwards, and indicate the average rise and fall of the BDI growth from the trough to the peak and vice versa and can serve as rough initial guides in our subsequent analysis.

The regularity of rise-and-fall for the BDI, in its levels and annual growth, that can be seen by the above identification of the turning points prompts us to consider the shape of its spectral density. For the rest of our discussion we concentrate solely on the annual growth series. Thus, in Figure 2 we plot (again) the series of the BDI annual growth and beneath it we plot the estimated spectral density of the series. There is one clear peak in the spectrum at frequency that corresponds to a period of 66 months or about 5.5 years (a smaller number can be seen before the 0.2 frequency mark corresponding to a 6 month period but is inconsequential). To gauge whether this dominant period is maybe a coincidence of the full sample we consider, in Figure 3, some sub-samples and their corresponding spectra. On the top left plot we have the series up to August, 2003 where the earlier behaviour can be seen along with the increase to higher levels towards 2004; on the top right plot we have the series from September, 2003 until May, 2008 where we can see the elongation of the cycle duration; in the left plot of the second row we have the series from June, 2008 until the end of the sample, in March, 2012. Here we have the spike in late 2009 and a complete absence of a regular cyclical pattern; in the right plot of the second row we have the series for the last two years where one can conjecture that we observe a re-emergence of some of the previous regularity of the cyclical pattern for the series. In the last two rows of the figure we present the estimated spectra for these sub-samples. For the first two sub-samples, i.e. ending in 2003 and 2008, we can clearly see the sharp peak at the same dominant frequency as in the one in Figure 2, i.e. for a period of about 5 years. These findings show that the period of the cycle in the BDI annual growth is quite close to that of the global business cycle (which, according to the literature ranges between 3 and 5 years). However, after the recent crisis and beyond 2008 the spectrum does not have any significant frequencies except close to the zero frequency (long-term trend), which corresponds to the post-crisis adjustment of the BDI. Finally, in the bottom right plot – where we have the spectrum for the last two years – we can see a dominant frequency at the low-frequency band that corresponds to a period of 3-years (the higher peaks at higher frequencies cannot be trusted here as they correspond to fast cycles which cannot be seen in this small sample). One can conjecture that this possibly signifies a post-crisis return to a more regular cyclical pattern, one that maybe will come to match again with the 5-year period we have seen that existed for most of the earlier part of BDI annual growth.

3.3 Coincidence and Synchronisation of the BDI Annual Growth

The presence of a, possibly regular, cyclical component is of interest but lacks further information that can be useful for decision making and forecasting. To do this we now turn to the statistics presented previously and see whether the cycles in BDI growth move together with those of other, related, variables.

In Table 2 we have some statistics on the coincidence and possible synchronisation of the annual change of the BDI with the annual change in a number of such variables. Two cycles are considered, an annual and a 5-year cycle and there are three measures of the degree of synchronisation: the mean difference, the correlation and the C-index (concordance index). Using the same notation as before, Y_t denotes the annual percentage change of the BDI and X_t stands consecutively for each of the following variables: COPPER, COTTON, TIN, MSCIEM, CRUDE, WHEAT, MSCIDEV, CORN, DXY, the exchange rate of GBP/USD and the SPREAD – all defined in the data Section. As mentioned earlier, the C-index is a practically useful measure as it allows us to see which variables have the strongest connection with the BDI. In Table 2 we have ranked the variables based on this measure.

We can see that the variables that have the highest coincidence with the BDI annual percentage change include the COPPER, COTTON and TIN in both the annual and the 5-year cycle - being in phase with the BDI more than 50% of the time in both cases. This is a finding that conforms with intuition, as these variables are commodities whose freight prices have feedback with the BDI itself. Then, the variables change based on the cycle length we consider. For the annual cycle we see that MSCIDEV, MSCIEM, CRUDE and WHEAT have a C-index of more than 50% while for the 5-year cycle the next important variables are the GBPUSD and MSCIDEV that have a C-index of more than 45%. Again, these variables conform to the underlying intuition of the factors that affect the BDI: we have the emerging and developed market indices, that can be thought as economic strength indicators which move in a procyclical fashion with the BDI; we then have two more commodities, one that is part of the regular cargo WHEAT, and another which affects the cost of transport CRUDE; finally we have the GBPUSD exchange rate which moves in relative concordance with the BDI.

However, measuring the concordance of the BDI with these other variables – while informative – is not sufficient. We are also interested in the statistical significance of cycle synchronisation. The z-statistics in the tables are for formally testing the hypotheses of positive synchronisation (z-statistic on mean difference) and of negative synchronisation (z-statistic on correlation). For the annual cycle, and for a 5% level of significance, we reject the hypothesis of positive synchronisation in favour of a negative one for the DXY and SPREAD variables. We thus have that the slope of the US Treasury curve (i.e. the 10-year minus 2-year yields) is also an important variable to watch, as the negative synchronisation implies that a steepening US yield curve is associated with higher BDI outcomes (one can think lower spreads as indicators of a "risk-on" global environment or of a time period of strong real growth). For the DXY index the result is not surprising as the under-performance of the US dollar against a basket of major currencies is associated with higher BDI outcomes as well (e.g. a cheaper dollar is associated with increased trade flows that raise the BDI from the demand side). For the 5-year cycle, and the same level of significance, we reject positive synchronisation for COTTON, TIN, SPREAD, CRUDE and MSCIEM and accept negative synchronisation for COTTON, SPREAD and CRUDE. We have asynchronicity for TIN and MSCIEM. Finally, CORN and WHEAT also appear to be asynchronous to BDI at this cycle length while GBPUSD, MSCIDEV, COPPER and TIN are positively synchronised. This last finding is of importance since it involves the GBPUSD exchange rate and the emerging markets index: the positive synchronisation of the pound is in line with the negative synchronisation of the DXY variable while the positive synchronisation of the emerging markets is intuitively explained by the demand-side associations of the growth path of the global economy.

3.4 How Does Cyclicality Affect Persistence?

With the use of a simple autoregressive model, we now ask whether the cyclical features of the BDI can explain the mean and the strength of the persistence of the BDI annual growth. A constant parameter autoregressive model might be OK for a forecasting comparison (which we perform later) but what if, a suitably adapted model, can show some structural light on the evolution of the BDI? Writing a functional autoregression as,

$$y_t = \phi_0(z_t) + \phi_1(z_t)y_{t-1} + \varepsilon_t \tag{6}$$

where z_t is a simple cosine of period of either 3 or 5 years. This model therefore links the state of the cycle, at or close to the dominant period, with the evolution of both the model mean $\mu(z_t) = \phi_0/\left[1 - \phi_1(z_t)\right]$ and the persistence parameter $\phi_1(z_t)$. The functional parameters are estimated by kernel methods as in Cai, Fan and Yao (2000).

The estimated functional coefficients are given in Figure 8. The panels include plots of the 3 and 5 years mean and autoregressive coefficient against the current state of the cycle which varies regularly over the interval [-1, 1], i.e. from the trough to the peak in the half-life of the cycle. Therefore the plots describe the cyclical variation of the coefficient estimates (one "reads" the horizontal axis both ways, from the beginning to the end and vice versa) that allows for a nice interpretation and shows the asymmetries present in the data.

Starting with the top two panels of the figure we can see that the 3-year and the 5-year cycles have different characteristics that are, however, both asymmetric with respect the current position in the cycle. They start off at the trough in the negative territory and in about the same value and then, as we move towards the peak, they increase but in a different fashion. The longer cycle means evolution is more smooth and has a marked increase as one goes towards the peak – which has a much larger amplitude than the corresponding value of the 3-year cycle. Furthermore, notice that the mean for the 5-year cycle is above zero during the first part of the recovery but not so during the second part (after the zero in the horizontal axis). This indicates that the initial uptrends after the troughs are possibly stronger and with higher acceleration. On the other hand, we have to be well into the recovery to see the exponential increase towards the peak of the cycle. What is interesting here is of course that we can now read the graph backwards: the way down is much faster, as we have an exponential decline, a modest correction and then we meet the trough before the cycle starts again. These results for the 5-year cycle are interesting because they refer to a period of variation quite close to the business cycle. The results for the 3-year cycle show more overall stability and higher variation closer to the position of the trough. We find that, at the end of the estimation sample of our data and for the 5-year cycle, we are almost near the trough (blue vertical line is close to but not yet -1) while for the 3-year cycle we are in about the middle of the cycle. The 5-year cycle positioning in the cycle tallies probably well with the results in Figure 3 where both indicate we might be going into an up-swing.

In the bottom two panels of Figure 8 we have the plots of the persistence (autoregressive) coefficient aligned with the position in the cycle. A high value on the vertical axis indicates higher persistence and thus it means that it is less likely to see a reversal in the series. The results here are quite clear for both the 5-year and the 3-year cycles: recovering from the trough is a period where the series has higher persistence than when is going down from a peak and there appears to be a threshold/critical point during the cycle where a break in persistence takes place, in the middle of the second part of the cycle.

All in all, we see that both the mean and the persistence of the BDI annual growth series exhibit changes that might be tied to the state of the BDI cycle, as captured by a simple cosine wave at (or close to) the dominant period found in the series spectrum.

4 Forecasting the BDI

4.1 Models

The potential presence of cyclicality in the BDI annual growth series, and the presence of variables that are pro- or counter-cyclical with the BDI, both suggest that they might be useful in forecasting the series into the future. Such an exercise goes beyond the relative ability of variables and models to produce (statistically) accurate forecasts and stretches into the realm of practical planning. We thus consider medium and long(er) term forecasts that go to six and twelve months ahead, horizons that are both practically useful and do not overtax the models that generate the forecasts. In such an exercise the choice of a benchmark is significant and we could have chosen among a variety of models. However, to ensure that any of the above findings does not bias the final ranking of the models we stick to the standard, a-theoretical, choice of an autoregressive model as the benchmark. We next turn to a presentation and justification of the rest of the models used in our forecasting exercise.

4.1.1 Trigonometric Regression

Using the information from the spectral analysis we start by considering the simplest model that can capture the cyclical patterns, i.e. a trigonometric regression. This model uses a combination of cosines or sines and cosines at pre-specified frequencies to explain the cyclical trends in the data. We use three frequencies for the fitting that correspond to periods of about 3, 4 and 5 year cycles. Then, two trigonometric models are described as,

$$\text{TRIG\#1} : y_t = \alpha + \sum_{j=1}^{3} \beta_j z_{tj} + \varepsilon_t,\ t = 1, \ldots, T, \tag{7}$$

$$\text{TRIG\#2} : y_t = \alpha + \sum_{j=1}^{3} (\beta_j z_{tj} + \gamma_j w_{tj}) + \varepsilon_t,\ t = 1, \ldots, T, \tag{8}$$

where $z_{tj} = \cos(2\pi\lambda_j t)$, $w_{tj} = \sin(2\pi\lambda_j t)$ are the transcendental factors evaluated at the three chosen frequencies denoted by λ_j. In Figure 6 we have the full sample fit from the trigonometric model when we fit the first model with cosines and the second model with both sines and cosines. The composite model explains 30% of the variability of the annual change of the BDI, a rather large number given its simplicity. In the top left panel of Figure 7 we have the sample fit for the period up to August, 2003 and the model now explains over 60% of the variability, a result that is in accordance to our previous discussion about the cyclical regularity during these years in the sample. Moving on the top right panel of the figure, where we have the sample fit for the period of September, 2003 to May, 2008, we can clearly see that the pattern of the larger cycles of 3 to 5 years tracks the series quite well with high explanatory power of over 80%. The fit is not suitable for the sub-sample in the bottom left panel where the data exhibits a large deviation from a regular cyclical pattern. However, in the bottom right panel of the figure where we have the data for the past two years, we can again see that the larger cycles track well the movement of the series and the explanatory power rises to 82%. These sample fits indicate that the second model is better than the first in capturing the cyclical variability in the BDI annual growth, hence this is the one we consider for the rest of the study and we denote it by TRIG in what follows.

4.1.2 Factor Selection via Principal Components

Next, we consider some models that are standard choices in the forecasting literature. One of the most widely used class of forecasting methods using variable reduction are factor methods. Factor methods have been at the forefront of developments in forecasting with large data sets and in fact started this literature with the influential work of Stock and Watson (2002a). The defining characteristic of most factor methods is that relatively few summaries of the large data sets are used in the forecasting equation, which thereby becomes a standard forecasting equation as it only involves a few variables. The assumption is that the co-movements across the indicator variables x_t, where $x_t = (x_{t1} \cdots x_{tN})'$ is a vector of dimension $N \times 1$, can be captured by a $r \times 1$ vector of unobserved factors $F_t = (F_{t1} \cdots F_{tr})'$, i.e.,

$$\tilde{x}_t = \Lambda' F_t + e_t, \tag{9}$$

where $\tilde{x}_t$ may be equal to x_t or may involve other variables such as, e.g., lags and leads of x_t and Λ is a $r \times N$ matrix of parameters describing how the individual indicator variables relate to each of the r factors, which we denote with the terms 'loadings'. In equation(9) e_t denotes a zero-mean $I(0)$ vector of errors that represents for each indicator variable the fraction of dynamics unexplained by F_t, the 'idiosyncratic components'. The number of factors is assumed to be small, meaning $r < \min(N, T)$. The main difference between different factor methods relates to how Λ is estimated.

The use of principal components (PC) for the estimation of factor models is, by far, the most popular factor extraction method. It has been popularised by Stock and Watson (2002a) and Stock and Watson (2002b), in the context of large data sets, although the idea had been well established in the traditional multivariate statistical literature. The method of principal components is simple. Estimates of Λ and the factors F_t are obtained by solving,

$$V(r) = \min_{\Lambda, F} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} (\tilde{x}_{ti} - \lambda_i' F_t)^2, \tag{10}$$

where λ_i is a $r \times 1$ vector of loadings that represent the N columns of $\Lambda = (\lambda_1 \cdots \lambda_N)$. One, non-unique, solution of equation (10) can be found by taking the eigenvectors corresponding to the r largest eigenvalues of the second moment matrix $\tilde{X}'\tilde{X}$, which then are assumed to represent the rows in Λ, and the resulting estimate of Λ provides the forecaster with an estimate of the r factors $\hat{F}_t = \hat{\Lambda}\tilde{x}_t$. To identify the factors up to a rotation, the variables are usually normalised to have zero mean and unit variance prior to the application of principal components; see Stock and Watson (2002a) and Bai (2003).

PC estimation of the factor structure is essentially a static exercise as no lags or leads of x_t are considered. One alternative is dynamic principal components, which, as a method of factor extraction, has been suggested in a series of papers by Forni, Hallin, Lippi and Reichlin (see, e.g., Forni et al. (2000) among others). Dynamic principal components are extracted in a similar fashion to static principal components but, instead of the second moment matrix, the spectral density matrices of the data at various frequencies are used. However, evidence suggests that static PC is a more effective and robust technique for forecasting and therefore this is our focus. We use the above method to extract the first principal component using the following indicator variables: CRUDE, MSCIDEV, MSCIEM, GBP/USD, DXY, COPPER, TIN, CORN, WHEAT, and COTTON.

It is important to notice here that, given the previous analysis, the cyclical effect is attempted to be captured using the explanatory variables due to their coincidental relationship with the BDI.

4.1.3 Linear Regressions & Broken Trends

The bulk of the forecasts is generated from linear models using the variables that capture the cyclical properties of the BDI series, as analysed in the previous section, plus combinations of these variables with the extracted PC factors. We thus consider the following generic regression model,

$$y_t = \alpha + \sum_{j=1}^{K} \beta_j x_{tj} + \sum_{i=1}^{2} \tau_t(\theta_i) + \varepsilon_t, \quad t = 1, \ldots, T, \tag{11}$$

where y_t denotes the annual growth of the BDI series, x_{tj} is the j-th explanatory variable for $j = 1, 2, ...K$ and $\tau_t(\theta_i)$ is a trend component explained below. We use the following sets of explanatory variables:

- Model PC: in this model $K = r$, where the first r factors from the use of PC of the previous section. We take r to be that number of factors that estimates at least 90% of the variance of all variables included in the PC analysis.
- Model COM (commodities): in this model $K = 6$ using the following commodities: CRUDE, COPPER, TIN, CORN, WHEAT and COTTON. These are the variables that remained after four rounds of top-down, sequential elimination – from the total number of explanatory variables – at declining levels of significance of 20%, 10%, 5% and 1%.
- Model CORN: in this model $K = 1$ and only the variable CORN is used.
- Model PC-trend: in this model $K = r$, with the r factors as above, plus a

number of broken quadratic trend components of the form,

$$\tau_t(\theta_i) = \sum_{s=1}^{2} (b_s t^s + c_s t^{s+1}) I(t > \theta_i) \tag{12}$$

where the thresholds θ_i are chosen to be the last 60 and 90 observations – so that a broken trend is estimated only for the latest portions of the data.

- Model COM-trend, where the trend components of the above equation are included in the commodities regressors.
- Model CORN-trend, where the trend components are included in the regression with CORN as a regressor.

As noted in the beginning of this section, our benchmark model is a simple autoregression that is described as,

$$y_t = \phi_0 + \phi_1 y_{t-1} + \varepsilon_t, \quad t = 1, \ldots, T, \tag{13}$$

As above, the cyclical effect is attempted to be captured using the explanatory variables due to their coincidental relationship with the BDI.

A final comment: the inclusion of monthly dummies did not lead to any significant improvement in forecasting performance and hence cases where these dummies were used are omitted from the presentation. However, full results are available to the interested reader on request by the authors.

4.1.4 Forecast Generation, Averaging and Evaluation

We perform a forecasting exercise using the projection method as described in Stock and Watson (2002a). This method, also known as the direct approach, is more robust in the presence of possible model mis-specification. The forecasts for any model m are then given by,

$$\widehat{y}_{t+h}^{f,m} = z_t' \widehat{\beta}^h, \tag{14}$$

where $\widehat{\beta}^h$ is obtained by regressing y_t on the lagged z_{t-h}, h denoting the forecast horizon.* z_t is an appropriately dimensioned vector of variables that come from either the trigonometric regression or the linear models described above.

We then specify the (rolling) estimation period R and the evaluation period P so that a summary of a standard pseudo out-of-sample forecasting algorithm is given as follows.

1. Use the rolling sample of R observations ($R = T - P - h$).

2. With any method described in this section obtain z_{t-h}, with $t = 1, 2, ..., R$.

3. Regress y_t on z_{t-h} and obtain $\widehat{\beta}^h$.

4. Calculate the forecasts of $\widehat{y}_{t+h}^{f,m}$ at periods $t = R+1, R+2, \ldots, R+h$ using sequentially the values of the explanatory variables $(z_{t-h+1}, z_{t-h+2}, \ldots, z_t)$ a period $t = R$ and the coefficient estimate $\widehat{\beta}^h$.

5. Repeat steps 2 to 4 by rolling the initial sample one period ahead, i.e. by setting $t = 2, 3, \ldots, R+1$ in step 2 and accordingly in steps 3 and 4.

Due to limited availability of data the number of rolling estimation periods has been set to $R = \{90, 180\}$.

Besides the direct forecasts that we obtain from the models we also consider various forecasts that come from model averaging. All averaging is done with equal model weights and we have, in the obvious notation:

1. $\widehat{y}_{t+h}^{f,A} = (\widehat{y}_{t+h}^{f,COM} + \widehat{y}_{t+h}^{f,CORN} + \widehat{y}_{t+h}^{f,PC})/3$, the average forecast of the models without a trend.

2. $\widehat{y}_{t+h}^{f,A-Trend} = (\widehat{y}_{t+h}^{f,COM-Trend} + \widehat{y}_{t+h}^{f,CORN-Trend} + \widehat{y}_{t+h}^{f,PC-Trend})/3$, the average forecast of the models with a trend.

*Since only lagged values are used as explanatory variables we do not have to face the endogeneity problem mentioned in the introduction. As a referee mentioned, we could still have an endogeneity problem if the regression error is autocorrelated and of the same order as the delay h used in the explanatory variables. However, our residual diagnostics do not suggest that our residuals suffer from autocorrelation and, therefore, this potential endogeneity source is not present.

3. $\widehat{y}_{t+h}^{f,A-Trig} = (\widehat{y}_{t+h}^{f,COM} + \widehat{y}_{t+h}^{f,CORN} + \widehat{y}_{t+h}^{f,PC} + \widehat{y}_{t+h}^{f,TRIG})/4$, the average forecast of the models without a trend plus the trigonometric regression.

4. $\widehat{y}_{t+h}^{f,A-Trig-Trend} = (\widehat{y}_{t+h}^{f,COM-Trend} + \widehat{y}_{t+h}^{f,CORN-Trend} + \widehat{y}_{t+h}^{f,PC-Trend} + \widehat{y}_{t+h}^{f,TRIG})/4$, the average forecast of the models with a trend plus the trigonometric regression.

5. $\widehat{y}_{t+h}^{f,A-Trig-Trend}$, the average of all seven forecasts.

At the end of this process we have gathered a total number of P forecast values for any horizon h from any model m. The forecast errors are then calculated as,

$$\widehat{e}_{t+h}^{f,m} = y_{t+h} - \widehat{y}_{t+h}^{f,m} \tag{15}$$

A final step in the forecast generation is an unbiasedness correction that we effect by adjusting the forecasts by the means of the (recursive) forecast errors. This is done so as to (smoothly) "correct" the forecasts as time elapses. We do this as follows,

$$\tilde{y}_{t+h}^{f,m} = \delta_h \widehat{y}_{t+h}^{f,m} + (1 - \delta_h)\frac{1}{t}\sum_{i=R}^{t-1} \widehat{e}_{i+h}^{f,m} \tag{16}$$

where δ_h is the smoothing factor. We use a sequence that, as h increases, progressively gives less weight to the forecast and more to the mean error factor, that is $\delta_h = [0.95 - 0.05(h-1)]$. After these smoothed forecasts are generated we have the new set of the forecast errors given by,

$$\tilde{e}_{t+h}^{f,m} = y_{t+h} - \tilde{y}_{t+h}^{f,m}. \tag{17}$$

Once these forecasts errors are available then evaluation statistics of interest can be computed. We are particularly interested in the Root Mean Squared Forecast Error (RMSFE) defined as,

$$RMSFE(h,m) = \frac{1}{P}\sum_{t=R}^{T-h}\left[\left(\tilde{e}_{t+h}^{f,m}\right)^2\right]^{\frac{1}{2}}. \tag{18}$$

We also calculate the Diebold and Mariano (1995) statistic for testing the predictive accuracy of different models. Here we use the two-sided test where the set of hypotheses is as follows:

- $H_0 : E[d_t] = 0$
- $H_A : E[d_t] \neq 0$,

where d_t is the loss differential defined as,

$$d_t = (e_{t+h}^{f,m_1})^2 - (e_{t+h}^{f,m_2})^2, \tag{19}$$

for two competing models m_1 and m_2; $m_1 \neq m_2$ and $t = R, .., T - h$. Then, the Diebold and Mariano (1995) test statistic is given by,

$$S = \frac{\overline{d}}{\left(\widehat{LRV}_{\overline{d}}/P\right)^{1/2}}, \tag{20}$$

with,

$$\overline{d} = \frac{1}{P}\sum_{t=R}^{T-h} d_t, LRV_{\overline{d}} = \gamma_0 + 2\sum_{j=1}^{\infty}\gamma_j, \gamma_j = cov\left(d_t, d_{t-j}\right), \tag{21}$$

where $LRV_{\overline{d}}$ is a consistent estimate of the long-run variance $\sqrt{P}\overline{d}$. Under the null of equal predictive accuracy, the statistic is distributed as $S \sim N(0, 1)$.

The Sign Success Ratio (SSR) is defined as the proportion of instances that the direction of the forecasts from each model is the same to the direction of the actual values and is given by,

$$SSR(h, m) = \frac{1}{P}\sum_{t=R+1}^{T-h} I\left[sgn(\Delta y_{t+h}) = sgn(\tilde{y}_{t+h}^{f,m} - y_{t+h})\right], \tag{22}$$

where $sgn(\bullet)$ denotes the sign operator and $I(\bullet)$ is an indicator variable which takes the value 1 if the signs are equal and 0 otherwise.

4.2 Forecasting Results & Discussion

In Tables 3 and 4 we report the relative (to the AR(1) benchmark) RMSFE, the p-value of the Diebold and Mariano (1995) test statistic and the detailed SSR of all models over the respective evaluation periods P. There are two evaluation periods that are dictated by our choice of rolling windows. The rolling window of 90 months allows us an out-of-sample evaluation period from 2001 to 2012 (118 months), while the rolling window of the 180 months allows us an out-of-sample evaluation period from 2009 to 2012 (27 months). In the first case we use the initial sample from 1993 to 2001 in order to estimate the forecasts and then we perform the rolling window exercise. The 2001 to 2012 is our out-of-sample cross-validation data. The choice of these evaluation periods is obvious, as the second period includes the post-Lehman collapse period that contains the bulk years of the financial crisis. In reading the tables a value greater than one indicates that the benchmark model is better while a value less than one indicates that the corresponding model is better.

Looking first at Table 3 one thing stands clear: short-horizon forecasts are no better than the benchmark, save possibly for the SSR of the CORN model and the Average (AVG) model. However, once we go beyond the one-month horizon the results are drastically different and, here, the usefulness of the explanatory variables and model averaging comes through. We can make two useful observations: (a) the RMSFE improvements are more evident in the 12-month rather than the 6-month horizon – and this horizon might be way more useful for practical planning. Although we still see that we reject the equal predictive accuracy at (at least) 90% level for PC, CORN, COM and CORN with trend, AVG and AVG-TRIG for $h = 6$. Note that the best model for both the 6-month and the 12-month horizon is the Average model that includes the trigonometric regression, with a close second the PC model (note that in the 12-month horizon the order is AVG-TRIG, CORN, AVG, PC but the differences are not large); there is, therefore, some useful predictive ability in the use of the chosen explanatory variables in longer forecasting horizons. Then, (b) the SSR improvements of the models are not uniform and not as pronounced as the RMSFE improvements. Here we see that the AVG model is the only consistent performer across all horizons, followed by the TRIG model and the COM model – these last two are better than the benchmark in both the 6- and 12-month horizons. These results, and the properties they imply for the models, are displayed in Figures 9 to 11. There we can clearly see how the benchmark's performance deteriorates with the forecasting horizon in terms of variance but not necessarily in terms of direction.

If we next turn to the results in Table 4 we see a qualitatively similar overview – supporting the results from the previous table. Looking at the RMSFE results, now we have that the AVG model is an overall consistent performer across all horizons which is, however, outperformed in the 6- and 12-month horizons by the simple TRIG model. This finding corroborates the earlier analysis on the cyclicality of BDI growth and suggests a continuation of the cyclical path of the earlier sample. Note that for the 6- and 12-month horizons the TRIG model has the best performance by far, even if in the 1-month horizon has the worst performance. The other models also perform very well in the 6- and 12-month horizons; notice, for example, that in the 12-month horizon we have the TRIG model, the AVG-TRIG with Trend model, the AVG-All model and the PC with Trend model, all outperforming the benchmark by a margin of at least 20% and a maximum of 40%. Turning next to the results on the SSR we see that the TRIG model is a consistent performer for the 1- and 6-month horizon but not the 12-month horizon. On the other hand, the AVG model now is the best overall performer in predicting the correct direction. It is still interesting to see which models dominate in predicting the direction for each horizon: we have the PC model to be the best in the 1-month ahead, the COM, TRIG, AVG and AVG with Trend models to tie for the best model in 6-months ahead and the PC and AVG-TRIG models to tie for the best model in 12-months ahead. A graphical illustration of the above results is given in Figures 12 to 14.

4.3 Model-based Investing & Risk Management

As noted in the introduction there is a clear need for adding another step in our analysis: even if we accept the presence of cyclicality in the data and, more so, 'believe' our forecasting models, one needs to see them put into a decision making context. In this section we attempt to do that by considering how these models can be put into use for investing in the BDI and/or performing risk management by utilising our forecast track record. If our approach in this section is successful then it opens up a practical use of the forecasting models and many other ideas can possibly be put to good use for anyone that is interested in the BDI path.

The idea here is very simple: if the signs of the forecasts are accurate (i.e. more accurate than a random sign choice) then we can invest or hedge the BDI by placing an appropriate 'bet', going either long or long/short depending on our risk preferences. If the model sign suggests a rise in BDI's annual return then we should be 'buying' the BDI and if the model sign suggests a fall in BDI's return then we should be ‘selling’ the BDI or, at least, avoiding exposure in the market. Alternatively, one can hedge the BDI by going (appropriately) long or short in any kind of asset that moves along with the BDI: for example, if the model sign suggests a rise in the BDI’s return and we want to cover (hedge) ourselves from a possible mistake then we should ‘buy’ the BDI and sell an asset that is positively correlated with the BDI (or buy an asset that is negative correlated with the BDI - the result is qualitatively the same).

Although the BDI is not directly tradable there are many ways in which one can track its path via tradable assets. For example, one can form a portfolio based on assets that are highly correlated with the BDI or consider future contracts. For illustrating the usefulness of the timing ability of the forecasting models we proceed as if the BDI was directly investable.

We next describe in some detail the way we conduct our investing experiment. We have, as noted before, two strategies: (i) a ’Long Only’ (L) and (ii) a ’Long/Short’ (LS). Both strategies are evaluated in the following manner:

1. We use a 90 months rolling window as our in-sample period and compute the 12-month ahead forecast for each of our forecasting models.

2. If the sign of the forecast is positive we open a new long ‘position’ on the BDI which we hold for the next 12 months; if the sign is negative we either stay out of the market (L) or open a new short ‘position’ on the BDI which we also hold for the next 12 months (LS).

3. We allow the window to roll one month ahead and we repeat the whole procedure; this implies that we are opening one new position each month which stays active for the next 12 months.

The performance of the positions thus obtained are to be compared with three benchmarks: one is the performance based on the forecasted signs of the AR(1) model (which is also our benchmark in the forecasting exercise described in the previous section), the other is the performance of just holding the BDI (and to mimic the above timing procedure we assume that we open a new long ‘position’ for the BDI every month) and the final is the time-series momentum of Moskowitz et al. (2012).

This last approach is particularly relevant as a benchmark, since it applies a sign-based methodology and it is implemented in a similar fashion with the proposed approach that we take. In particular, in the time series momentum one looks at the past sign of a series of returns of an asset and goes long or long/short based on it. There is, therefore, a similarity but also a crucial difference between the momentum approach and ours: in the former the past is used and is believed that its sign is propagated into the future while in the latter a model-based sign forecast is used.

If there is cyclicality that is being captured by the forecasted signs then our suggested procedure should be able to capture it: when the BDI is forecasted to fall over the next 12-months and we either eliminate our exposure into it or even go against it then we should do better by just holding on to it. Furthermore, since we are using a long-term forecast that goes into the next year we should have the trigonometric models perform better than other ones, including all three benchmarks. Our results are given in Tables 5 and 6 and also in Figures 15 and 16. In the tables we present various statistics on the investment performance of the suggested approach, the difference between the two is the evaluation period – in the second of these tables the evaluation starts one year before the last financial crisis.

The main result that can be seen immediately from both tables is that there is economic value in the use of model sign-based timing. In particular, the predicted signs that are based on the trigonometric model, and the average of models that include the trigonometric component, have the best economic performance: they have the highest cumulative return, the highest Sharpe ratios and the lowest maximum drawdowns.†

In Table 5 the sign-based performance of the pure trigonometric model is by far the best by any measure. For the long-only approach the annualised Sharpe ratio exceeds 2, which compares to a 1.32 value for the BDI, a 0.49 for the AR(1) model and a 0.67 for the time series momentum. It thus outperforms the three benchmarks by a wide margin in terms of risk-adjusted returns. Furthermore, the trigonometric model has the second lowest maximum drawdown (the lowest is from the average of models that include the trigonometric component) and it provides an additional 20% cumulative return compared to just holding on to the BDI. The results become even better when we consider the long-short approach, clearly indicating that there

†Values for average, volatility and Sharpe ratio are annualised.

are indeed alternating signs in the future path of the BDI returns which can be exploited via the forecasting models. Furthermore, notice that not only the trigonometric model but practically all forecasting models that we suggest outperform the benchmarks: see the PC model and the CORN model in particular. These results strongly suggest not only that the forecasting exercise of the previous section was not futile but, on the contrary, it provides us with tools to be exploited in a BDI risk-management context.

To further pursue the potential of what is presented above consider Figure 15 and notice how, with the onset of the financial crisis in 2008 the return of holding into BDI is falling until 2009. What if we performed our analysis starting a year before the crisis in 2007? Would we still have been able to reduce our exposure? Although the answer is affirmative and can be seen from 15, we repeat the analysis and evaluation statistics and present the results in Table 6 and Figure 16. As can be seen clearly from the figure, the steep fall of the BDI during the crisis is just avoided - either by staying out of being exposed to the BDI (top panel, long-only) or by going against it (bottom panel, long/short). The statistics at Table 6 just re-affirm our previous results on the value of sign-based timing.

All in all, the results of this section support our earlier findings and, moreover, they transform them into practical tools that can be used from anyone who wishes to manage exposure to the future path of the BDI. There are caveats, obviously, to what we just presented (such as that the BDI is not directly tradable) but the overall good performance of the strategies which are based on model forecasts is such that leaves room for many different ways for further improvements.

5 Conclusions

Our overall analysis provides several interesting and novel results about the evolution of BDI annual growth. First, the contribution of the paper to the literature is the cyclical analysis of the series at different levels. Past research was limited to seasonal analysis (see Kavussanos and Alizadeh (2001), Kavussanos and Alizadeh (2002)). We find that there is a strong cyclical pattern of cycle duration of between 3 to 5 years and that this pattern is relatively stable across time. We then examine the cycle synchronisation of the BDI with other variables and we find that the BDI is (a) negatively synchronised with a weaker dollar and the spread of the US Treasuries and (b) positively synchronised with the GBPUSD exchange rate, the emerging market equities, copper and tin. The mean and persistence of the BDI annual growth can also be linked to the state of the cycle and we provide evidence on this via a functional autoregression, which allows us a nice visual interpretation on this link.

Second, we perform a comprehensive forecasting performance evaluation by considering a variety of models and model averages that incorporate explanatory variables – carefully selected by top-down elimination – and the cyclical component found in the first part of our analysis. The results of our forecasting exercise show that performance gains are possible when using auxiliary information, either in the form of explanatory variables or in the form of the cyclical component of the BDI. These gains are not uniform across all models examined and are concentrated mainly in the medium- and longer-term forecasting horizons. However, in the cases where outperformance of the benchmark is found we can see several occasions that this is of a rather large magnitude. A judicious choice of models, that incorporate the features that affect the BDI, can thus lead to good forecasting performance and aid in planning and management in using the future direction of the BDI.

Finding cyclicality in economic time series might be considered old-fashioned but in the present case we cannot refute it easily. Not only the models we present exhibit very good statistical forecasting performance, they can also be used for controlling financial exposure and risk to the BDI. In the last part of our analysis we perform a risk-management experiment where 12-month ahead forecasts are used to decide whether or not to invest in the BDI. Within the limitations we discussed above, the results on this third part of our analysis strongly support the long(er) term potential benefits of using sign-based timing for investing or hedging the BDI. Not only do we find that the trigonometric model gives the best economic performance in this experiment we also find that all of our forecasting models provide a better decision-making tool than any of the three benchmarks we employ.

Our results now open up a very interesting avenue of future research: how can we construct a realisable risk management system that will utilise the model signs and associated information so that one can exercise a higher degree of control when exposed to BDI fluctuations. Such a system will be used for both investing in and hedging BDI risk and should depend on assets that are immediately available for trading. We are currently pursuing this lien of research.

References

Adland, R., Cullinane, K., 2006. The non-linear dynamics of spot freight rates in tanker markets. Transportation Research Part E 42, 211-224.

Bai, J., 2003. Inferential theory for factor models of large dimension. Econometrica 71, 135-173.

Batchelor, R., Alizadeh, A. H., Visvikis, I., 2007. Forecasting sport and forward prices in the international freight market. International Journal of Forecasting 23, 101-114.

Beenstock, M., Vergottis, A., 1989a. An econometric model of the world shipping market for dry cargo, freight and shipping. Applied Economics 21, 339-356.

Beenstock, M., Vergottis, A., 1989b. An econometric model of the world tanker market. Journal of Transport Economics and Policy 23(3), 236-280.

Bornozis, N., 2006. Dry bulk shipping: the engine of global trade. Barron's.

Cai, Z., Fan J., Yao, Q., 2000. Functional-coefficient models for nonlinear time series. Journal of the American Statistical Association 95(451), 941-956.

Cullinane, K. P. B., Mason, K. J., Cape, M. B., 1999. Forecasting the baltic freight index: Box-Jenkins revisited. in Meersman, H., Van de Voorde, E., Winkelmans, W. (Eds.), World Transport Research, Selected Proceedings of the 8th World Conference on Transport Research Volume I: Transport Modes and Systems, Elsevier Science, Amsterdam, 73- 86.

Denning, K. C., Riley, W. B., DeLooze, J. P., 1994. Baltic freight futures: Random walk or seasonally predictable? International Review of Economics & Finance, 3(4), 399-428.

Diebold, F. X., Mariano, R. S., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics, 13, 253-263.

Driehuis, W., 1970. An econometric analysis of liner freight rates. Weltwirtschaftliches Archiv 104, 96-119.

Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2000. The generalised factor model: identification and estimation. Review of Economics and Statistics (82), 540-554.

Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2005. The generalised factor model: one-sided estimation and forecasting. Journal of the American Statistical Association 100(471), 830-840.

Giannakoulis, F., Slorer, O., 2012. Global shipping: time to buy containerships? Morgan Stanley Research, North America.

Goulas, L., Skiadopoulos, G., 2012. Are freight futures markets efficient? Evidence from IMAREX. MInternational Journal of Forecasting 28(3), 644-659.

Harding, D., 2008. Detecting and forecasting business cycle turning points. MPRA Paper No. 33583.

Harding, D., Pagan, A., 2006. Synchronization of cycles. Journal of Econometrics 132, 59-79.

Kavussanos, M. G., 1996. Comparisons of volatility in the dry-cargo ship sector: spot versus time charters. Journal of Transport Economics and Policy 30(1), 67-82.

Kavussanos, M. G., Alizadeh, A. H., 2001. Seasonality patterns in dry bulk shipping spot and time charter freight rates. Transportation Research Part E 37, 443-467.

Kavussanos, M. G., Alizadeh, A. H., 2002. Seasonality patterns in tanker spot freight rates. Economic Modelling 19, 747-782.

Koekebakker, S., Adland, R., Sodal, S., 2006. Are spot freight rates stationary? Journal of Transport Economics and Policy 40(3), 449-472.

Lin, F., Sim, N. C. S., 2013. Trade, income and the Baltic Dry Index. European Economic Review, 59, 1-18.

Marlow, P. B., Gardner, B., 1980. Thoughts on the dry bulk shipping sector. Journal of Industrial Economics 29(1), 71-84.

Moskowitz, T. J., Ooi, Y. H., Pedersen, L. H., 2012. Time series momentum. Journal of Financial Economics 104, 228-250.

Nomura Equity Research, 2012. Anchor report. 02-2012, 1-44.

Stock, J., Watson, M., 2002a. Forecasting using principal components from a large number of predictors. Journal of the American Statistical Association 297, 1167-1179.

Stock, J., Watson, M., 2002b. Macroeconomic forecasting using diffusion indexes. Journal of Business & Economics Statistics 20, 147-162.

6 Tables

	BDI	CRUDE	MSCIDEV	MSCIEM	GBPUSD	DXY	COPPER	TIN	CORN	WHEAT	COTTON	SPREAD
Mean	0.1989	0.1369	0.0661	0.1088	-0.0003	-0.0028	0.1258	0.1183	0.1001	0.0700	0.0843	1.1260
Median	0.0053	0.0894	0.1128	0.1223	0.0048	-0.0025	0.0569	0.0285	0.0251	0.0166	0.0271	1.0240
Std.Dev.	0.7197	0.3461	0.1773	0.3053	0.0892	0.0838	0.3703	0.3271	0.3308	0.3000	0.3462	0.9457
Skewness	1.7650	0.6330	-0.8601	0.1018	-0.9020	0.1355	1.0897	0.7361	1.0297	1.1405	1.1781	0.1697
Kurtosis	8.6970	3.8140	3.7632	2.6505	4.0599	2.2179	4.6667	2.6693	3.7788	4.7330	5.5186	1.5977

	BDI	CRUDE	MSCIDEV	MSCIEM	GBPUSD	DXY	COPPER	TIN	CORN	WHEAT	COTTON	SPREAD
BDI	1.0000											
CRUDE	0.3539	1.0000										
MSCIDEV	0.3702	0.3012	1.0000									
MSCIEM	0.5054	0.5010	0.6931	1.0000								
GBPUSD	0.4900	0.3118	0.4415	0.3811	1.0000							
DXY	-0.5070	-0.3161	-0.1751	-0.2740	-0.7270	1.0000						
COPPER	0.4626	0.4857	0.4652	0.6021	0.4885	-0.4447	1.0000					
TIN	0.4815	0.4219	0.4372	0.4671	0.5838	-0.5961	0.5679	1.0000				
CORN	0.0623	0.1577	0.1686	0.2153	0.2381	-0.3346	0.1451	0.5318	1.0000			
WHEAT	0.2560	0.2625	0.1507	0.3233	0.3379	-0.4631	0.2892	0.5925	0.7173	1.0000		
COTTON	0.3019	0.2268	0.2661	0.3365	0.2924	-0.4142	0.4056	0.5367	0.4084	0.4443	1.0000	
SPREAD	0.0884	-0.0714	-0.2532	0.0698	-0.0173	-0.2838	0.0182	0.1753	0.1028	0.0279	0.3758	1.0000

Table 1: Descriptive statistics and correlations of annual growth rates of all variables.

	Annual Cycle			5-years Cycle		
	Mean Difference	Correlation	Concordance	Mean Difference	Correlation	Concordance
COPPER	0.018	0.463	0.662	-0.164	0.302	0.593
(z-statistic)	0.38	4.29		-2.69	1.59	
COTTON	0.000	0.302	0.602	0.181	0.482	0.554
(z-statistic)	0.00	2.62		2.38	4.05	
TIN	0.050	0.482	0.597	0.076	0.463	0.494
(z-statistic)	0.73	2.96		0.60	2.07	
MSCIEM	0.046	0.505	0.593	0.152	0.490	0.472
(z-statistic)	0.82	3.32		1.34	2.58	
CRUDE	-0.055	0.354	0.558	0.193	0.370	0.450
(z-statistic)	-0.95	2.07		1.40	1.96	
WHEAT	-0.037	0.256	0.541	0.099	0.062	0.398
(z-statistic)	-0.52	1.69		0.91	0.41	
MSCIDEV	0.142	0.370	0.519	0.263	0.088	0.381
(z-statistic)	1.42	2.13		2.09	0.44	
CORN	-0.005	0.062	0.476	0.082	0.256	0.359
(z-statistic)	-0.07	0.48		0.67	1.68	
DXY	0.183	-0.507	0.463	0.275	0.354	0.329
(z-statistic)	2.18	-2.14		2.05	1.48	
GBP/USD	0.160	0.490	0.433	0.567	0.505	0.320
(z-statistic)	1.90	2.03		8.43	3.97	
SPREAD	0.178	0.088	0.424	0.216	-0.507	0.234
(z-statistic)	2.16	0.53		1.39	-3.78	

Table 2: Cyclical analysis and synchronicities of the annual percentage change of the BDI related to other variables.

Notes. Entries are the estimates and their z-statistics for the mean differences $\mathsf{E}\left[S_{ty} - S_{tx}\right]$, the correlation ρ_S and the concordance index I. The z-statistics are based in GMM standard errors with automatic lag selection. The null hypothesis for strong positive synchronisation corresponds to $\mathsf{E}\left[S_{ty} - S_{tx}\right] = 0$ and the null hypothesis of strong negative synchronisation corresponds to $\rho_S = 0$. $\mathbf{y}$ denotes the 12-month change of the BDI and $\mathbf{x}$ denotes each of the variables in the Table.

Rolling Window: 90, Evaluation Period: 2001 - 2012									
	Relative RMSFE			DM p-value			SSR		
Model	1-Step	6-Steps	12-Steps	1-Step	6-Steps	12-Steps	1-Step	6-Steps	12-Steps
PC	1.002	0.812	0.650	0.969	0.085	0.001	53.8%	52.1%	48.7%
COM	1.161	0.904	0.661	0.015	0.377	0.003	53.8%	53.8%	52.1%
CORN	1.008	0.829	0.628	0.739	0.014	0.000	60.7%	47.9%	54.7%
PC-Trend	1.331	1.161	0.809	0.000	0.143	0.184	46.2%	42.7%	48.7%
COM-Trend	1.468	1.248	0.862	0.000	0.045	0.322	42.7%	46.2%	50.4%
CORN-Trend	1.178	1.048	0.758	0.013	0.606	0.013	44.4%	43.6%	48.7%
TRIG	1.648	0.795	0.578	0.001	0.110	0.000	48.7%	52.1%	54.7%
AVG	1.013	0.824	0.638	0.738	0.052	0.000	57.3%	50.4%	51.3%
AVG-Trend	1.262	1.087	0.771	0.001	0.361	0.047	41.9%	44.4%	51.3%
AVG-TRIG	1.057	0.787	0.615	0.430	0.040	0.000	54.7%	54.7%	50.4%
AVG-TRIG-Trend	1.216	0.964	0.700	0.028	0.677	0.010	45.3%	45.3%	50.4%
AVG-ALL	1.105	0.894	0.665	0.097	0.193	0.003	48.7%	47.9%	49.6%
AR(1)	1	1	1	1	1	1	55.6%	49.6%	51.3%

Table 3: Forecasting exercise for the annual percentage change of the BDI. Reporting averages over 118 evaluation months

Notes. RMSFE denotes the relative root mean squared forecast error of each method to the benchmark. DM denotes the p-value of the two-sided Diebold and Mariano (1995) statistic using the squared difference of the forecast error of each method relative to the benchmark. SSR denotes the sign success ratio of each method. PC denotes the principal components method, COM denotes the linear regression model using CRUDE, COPPER, TIN, CORN, WHEAT and COTTON explanatory variables, CORN denotes the linear regression model using CORN variable, Trend denotes the use of the previously mentioned models including a deterministic trend estimate, AVG denotes the model average of the PC, COM and CORN models, AVG-Trend denotes the model average of the PC-Trend, COM-Trend and CORN-Trend models, AVG-TRIG denotes the model average of the PC, COM, CORN and TRIG models, AVG-TRIG-Trend denotes the model average of the PC-Trend, COM-Trend, CORN-Trend and TRIG models and AVG-ALL denotes the model average of all models. The benchmark model is the AR(1).

Rolling Window: 180, Evaluation Period: 2009 - 2012									
	Relative RMSFE			DM p-value			SSR		
Model	1-Step	6-Steps	12-Steps	1-Step	6-Steps	12-Steps	1-Step	6-Steps	12-Steps
PC	0.986	0.834	0.895	0.675	0.105	0.519	63.0%	48.1%	59.3%
COM	0.990	0.853	0.824	0.748	0.129	0.039	48.1%	51.9%	48.1%
CORN	1.002	0.883	0.928	0.902	0.147	0.670	48.1%	40.7%	51.9%
PC-Trend	1.140	1.012	0.793	0.064	0.894	0.104	48.1%	40.7%	51.9%
COM-Trend	1.172	1.086	0.831	0.057	0.681	0.000	48.1%	44.4%	37.0%
CORN-Trend	1.065	0.951	0.905	0.258	0.522	0.435	55.6%	40.7%	48.1%
TRIG	1.561	0.724	0.574	0.091	0.210	0.062	55.6%	51.9%	51.9%
AVG	0.986	0.853	0.880	0.568	0.106	0.371	59.3%	51.9%	55.6%
AVG-Trend	1.101	0.963	0.631	0.155	0.687	0.180	48.1%	51.9%	44.4%
AVG-TRIG	1.007	0.798	0.792	0.951	0.140	0.000	55.6%	48.1%	59.3%
AVG-TRIG-Trend	1.073	0.871	0.569	0.567	0.263	0.112	55.6%	48.1%	48.1%
AVG-ALL	1.018	0.853	0.675	0.820	0.146	0.008	59.3%	44.4%	51.9%
AR(1)	1	1	1	1	1	1	51.9%	48.1%	55.6%

Table 4: Forecasting exercise for the annual percentage change of the BDI. Reporting averages over 27 evaluation months

Notes. RMSFE denotes the relative root mean squared forecast error of each method to the benchmark. DM denotes the p-value of the two-sided Diebold and Mariano (1995) statistic using the squared difference of the forecast error of each method relative to the benchmark. SSR denotes the relative sign success ratio of each method to the benchmark. PC denotes the principal components method, COM denotes the linear regression model using CRUDE, COPPER, TIN, CORN, WHEAT and COTTON explanatory variables, CORN denotes the linear regression model using CORN variable, Trend denotes the use of the previously mentioned models including a deterministic trend estimate, AVG denotes the model average of the PC, COM and CORN models, AVG-Trend denotes the model average of the PC-Trend, COM-Trend and CORN-Trend models, AVG-TRIG denotes the model average of the PC, COM, CORN and TRIG models, AVG-TRIG-Trend denotes the model average of the PC-Trend, COM-Trend, CORN-Trend and TRIG models and AVG-ALL denotes the model average of all models. The benchmark model is the AR(1).

	Rolling Window: 90, Evaluation Period 2001-2012									
	Long Only					Long/Short				
	Average	Volatility	Sharpe	Cumulative	Drawdown	Average	Volatility	Sharpe	Cumulative	Drawdown
PC-Trend	0.041	0.029	1.402	0.492	0.089	0.040	0.032	1.246	0.476	0.089
CORN-Trend	0.039	0.028	1.428	0.468	0.058	0.037	0.033	1.128	0.427	0.085
COM-Trend	0.027	0.029	0.936	0.297	0.092	0.012	0.034	0.337	0.113	0.219
PC	0.052	0.028	1.865	0.656	0.040	0.061	0.029	2.094	0.819	0.026
CORN	0.054	0.027	1.982	0.686	0.025	0.065	0.029	2.265	0.883	0.030
COM	0.048	0.027	1.733	0.589	0.047	0.053	0.031	1.726	0.673	0.042
TRIG	0.055	0.027	2.045	0.710	0.038	0.068	0.028	2.414	0.937	0.038
AVG-Trend	0.036	0.030	1.197	0.413	0.092	0.029	0.033	0.875	0.324	0.112
AVG	0.050	0.028	1.799	0.631	0.045	0.058	0.030	1.951	0.763	0.035
AVG-TRIG-Trend	0.041	0.029	1.424	0.489	0.089	0.040	0.032	1.229	0.469	0.110
AVG-Trig	0.054	0.027	2.009	0.698	0.023	0.066	0.028	2.340	0.911	0.026
AVG-ALL	0.046	0.028	1.634	0.572	0.069	0.051	0.031	1.638	0.638	0.071
AR(1)	0.012	0.024	0.491	0.121	0.149	-0.018	0.034	-0.546	-0.171	0.254
TSM	0.015	0.023	0.668	0.160	0.124	-0.011	0.034	-0.336	-0.112	0.223
BDI	0.042	0.032	1.318	0.506	0.092	0.042	0.032	1.318	0.506	0.092

Table 5: Investing in the BDI. Reporting annualised statistics over 118 evaluation months

Notes. Average denotes the annualised mean of returns. Volatility denotes the annualised standard deviation of returns. Sharpe Ratio denotes the annualised average/standard deviation ratio of returns. Cumulative denotes the cumulative return. Drawdown denotes the maximum drawdown of the cumulative return. PC denotes the strategy that invests in the BDI using the signs obtained from the principal components method, COM denotes the strategy that invests in the BDI using the signs obtained from the linear regression model using CRUDE, COPPER, TIN, CORN, WHEAT and COTTON explanatory variables, CORN denotes the strategy that invests in the BDI using the signs obtained from the linear regression model using CORN variable, Trend denotes the strategy that invests in the BDI using the signs obtained from previously mentioned models including a deterministic trend estimate, AVG denotes the strategy that invests in the BDI using the model average of the PC, COM and CORN models, AVG-Trend denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC-Trend, COM-Trend and CORN-Trend models, AVG-TRIG denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC, COM, CORN and TRIG models, AVG-TRIG-Trend denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC-Trend, COM-Trend, CORN-Trend and TRIG models and AVG-ALL denotes the strategy that invests in the BDI using the signs obtained from the model average of all models. The benchmarks are the strategies that invest in the BDI using the signs obtained from the AR(1) and Time Series Momentum (TSM) models. BDI denotes the actual percentage change of the BDI series.

	Rolling Window: 90, Evaluation Period 2007-2012									
	Long Only					Long/Short				
	Average	Volatility	Sharpe	Cumulative	Drawdown	Average	Volatility	Sharpe	Cumulative	Drawdown
PC-Trend	0.028	0.033	0.848	0.154	0.089	0.026	0.036	0.724	0.144	0.089
CORN-Trend	0.028	0.030	0.949	0.157	0.058	0.027	0.036	0.747	0.149	0.059
COM-Trend	0.012	0.031	0.384	0.062	0.092	-0.005	0.037	-0.143	-0.031	0.183
PC	0.045	0.031	1.466	0.263	0.040	0.061	0.033	1.859	0.371	0.026
CORN	0.046	0.029	1.546	0.267	0.025	0.062	0.032	1.908	0.380	0.030
COM	0.033	0.030	1.102	0.187	0.047	0.037	0.036	1.038	0.210	0.042
TRIG	0.048	0.030	1.602	0.282	0.038	0.066	0.032	2.093	0.412	0.038
AVG-Trend	0.024	0.032	0.728	0.129	0.092	0.018	0.037	0.491	0.096	0.092
AVG	0.040	0.031	1.313	0.233	0.045	0.052	0.034	1.522	0.307	0.035
AVG-TRIG-Trend	0.026	0.032	0.827	0.144	0.089	0.023	0.037	0.629	0.124	0.110
AVG-Trig	0.047	0.030	1.596	0.278	0.023	0.065	0.032	2.049	0.405	0.026
AVG-ALL	0.036	0.032	1.128	0.202	0.069	0.042	0.035	1.195	0.242	0.071
AR(1)	-0.004	0.022	-0.189	-0.023	0.149	-0.038	0.036	-1.060	-0.182	0.254
TSM	-0.001	0.021	-0.070	-0.009	0.124	-0.032	0.036	-0.894	-0.158	0.223
BDI	0.029	0.036	0.808	0.162	0.092	0.029	0.036	0.808	0.162	0.092

Table 6: Investing in the BDI. Reporting annualised statistics over 63 evaluation months

Notes. Average denotes the annualised mean of returns. Volatility denotes the annualised standard deviation of returns. Sharpe Ratio denotes the annualised average/standard deviation ratio of returns. Cumulative denotes the cumulative return. Drawdown denotes the maximum drawdown of the cumulative return. PC denotes the strategy that invests in the BDI using the signs obtained from the principal components method, COM denotes the strategy that invests in the BDI using the signs obtained from the linear regression model using CRUDE, COPPER, TIN, CORN, WHEAT and COTTON explanatory variables, CORN denotes the strategy that invests in the BDI using the signs obtained from the linear regression model using CORN variable, Trend denotes the strategy that invests in the BDI using the signs obtained from previously mentioned models including a deterministic trend estimate, AVG denotes the strategy that invests in the BDI using the model average of the PC, COM and CORN models, AVG-Trend denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC-Trend, COM-Trend and CORN-Trend models, AVG-TRIG denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC, COM, CORN and TRIG models, AVG-TRIG-Trend denotes the strategy that invests in the BDI using the signs obtained from the model average of the PC-Trend, COM-Trend, CORN-Trend and TRIG models and AVG-ALL denotes the strategy that invests in the BDI using the signs obtained from the model average of all models. The benchmarks are the strategies that invest in the BDI using the signs obtained from the AR(1) and Time Series Momentum (TSM) models. BDI denotes the actual percentage change of the BDI series.

7 Figures

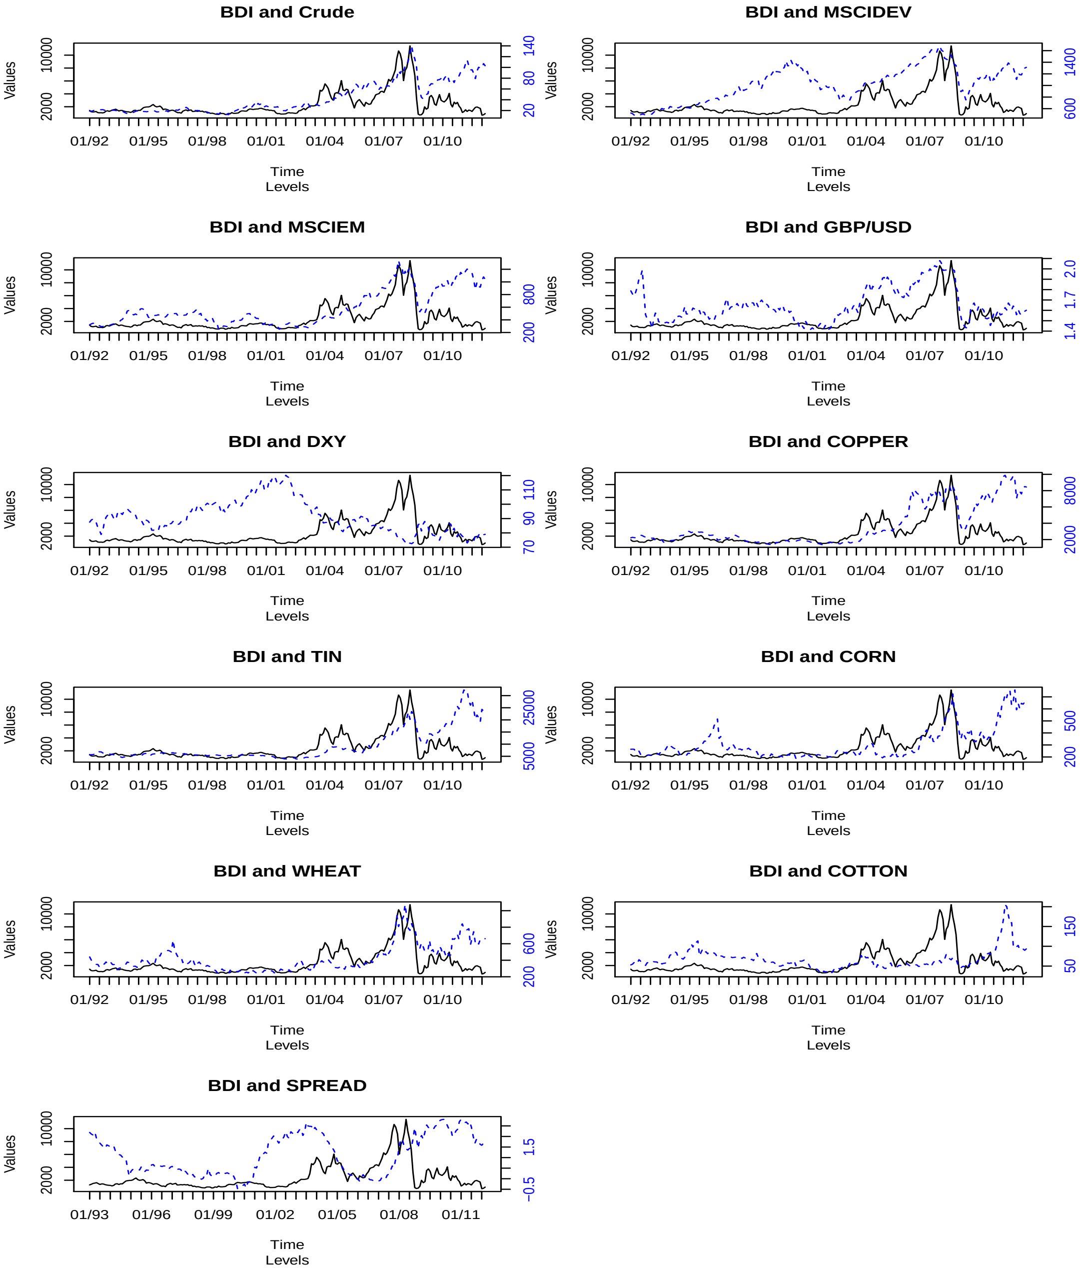

Figure 1: Comparing BDI to other variables using levels.

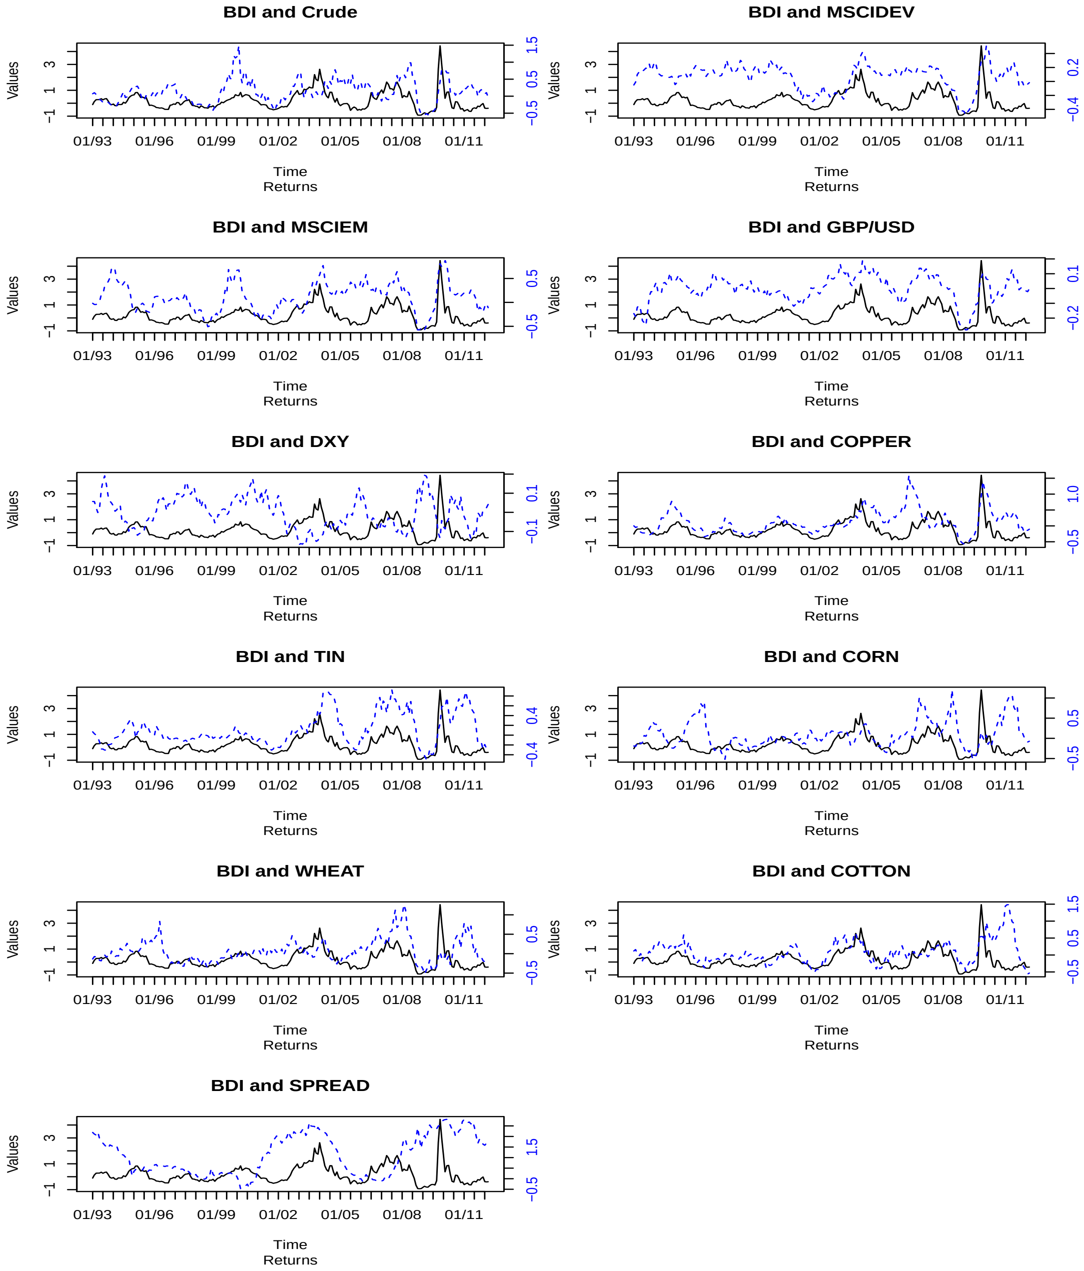

Figure 2: Comparing BDI to other variables using the annual percentage change.

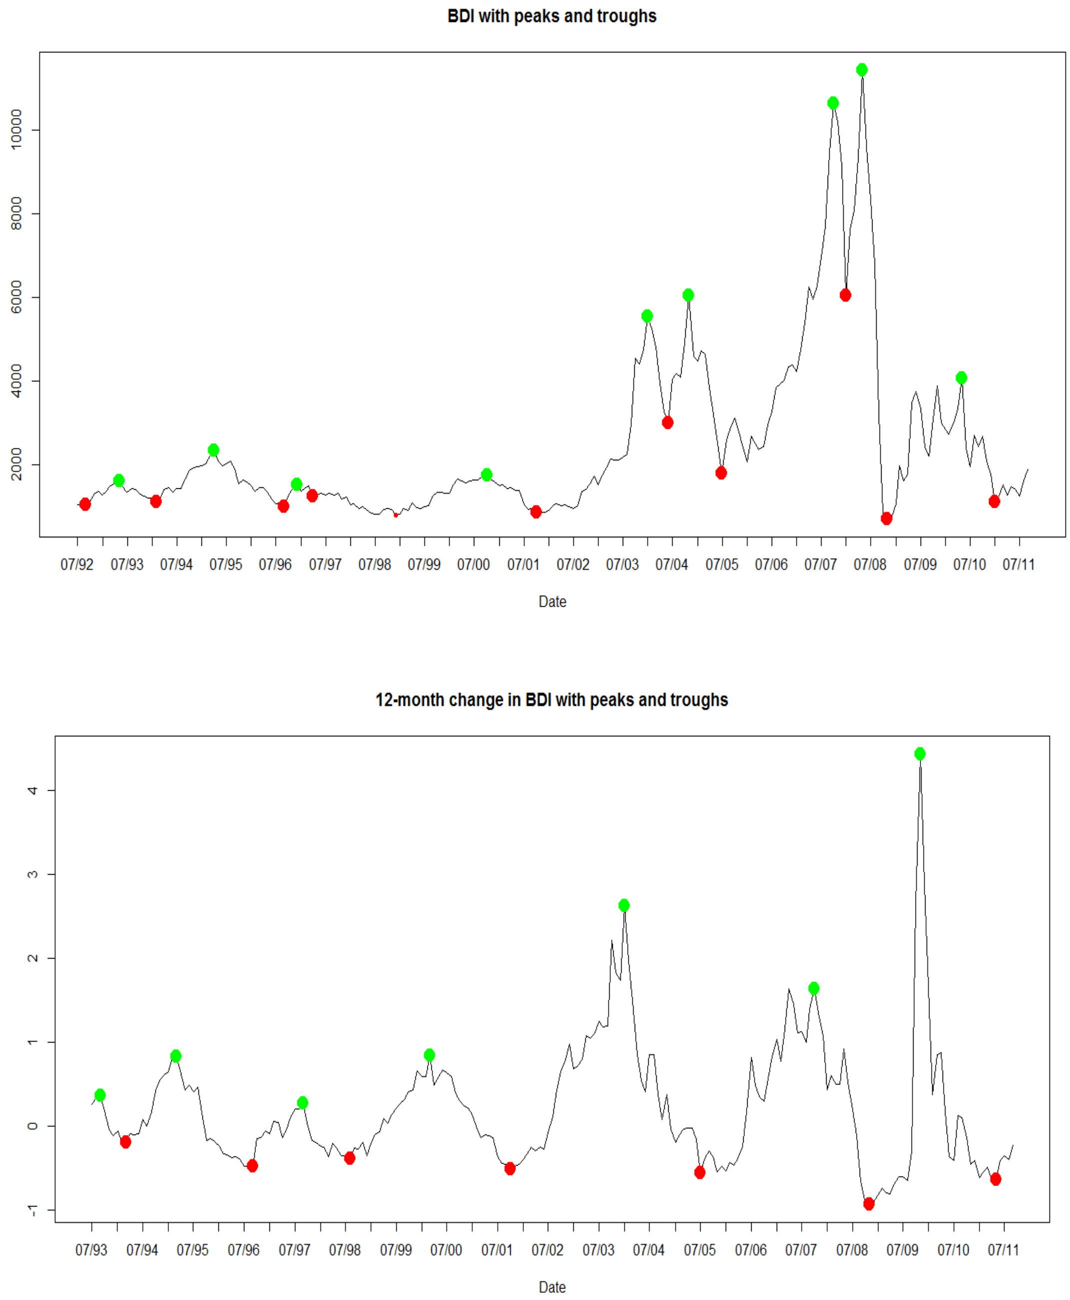

Figure 3: BDI series and annual percentage change with peaks (green) and troughs (red).

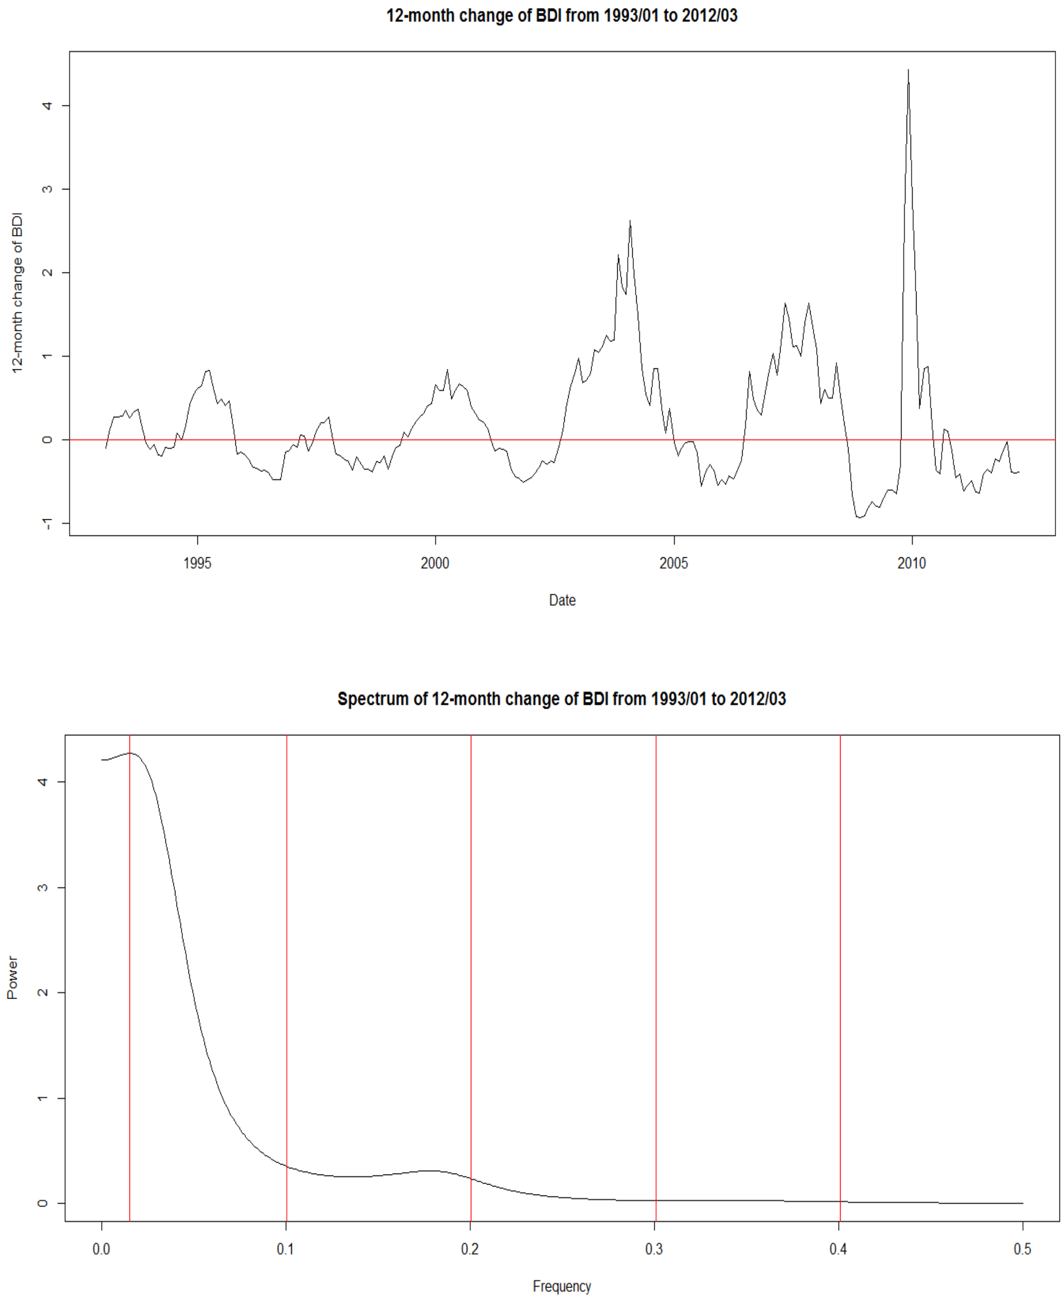

Figure 4: annual percentage change of the BDI and its spectrum.

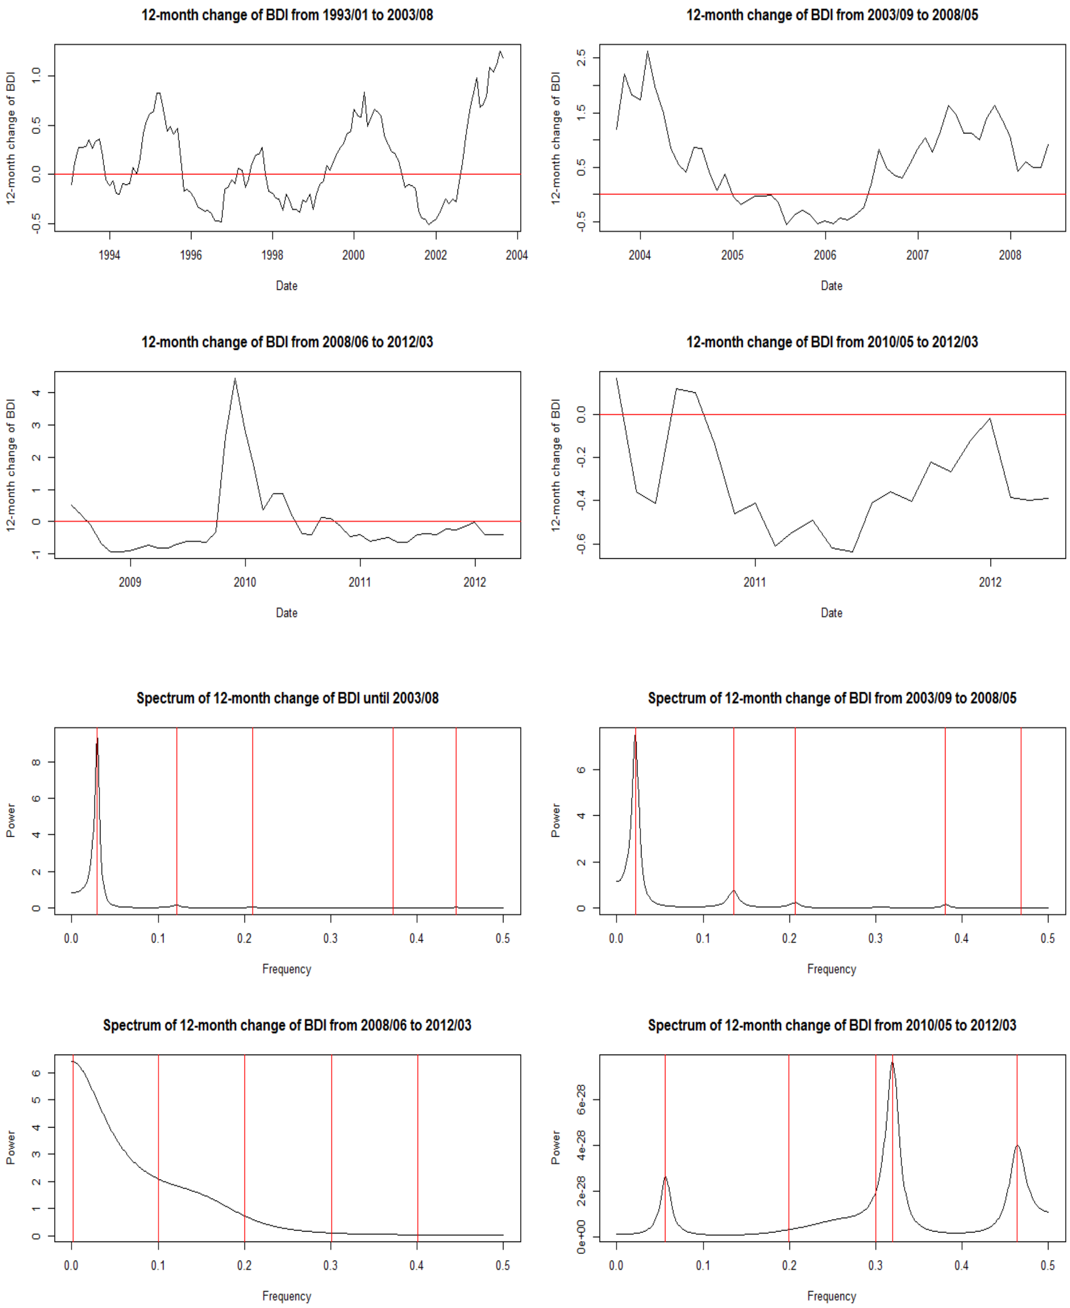

Figure 5: annual percentage change of the BDI and its spectrum across periods.

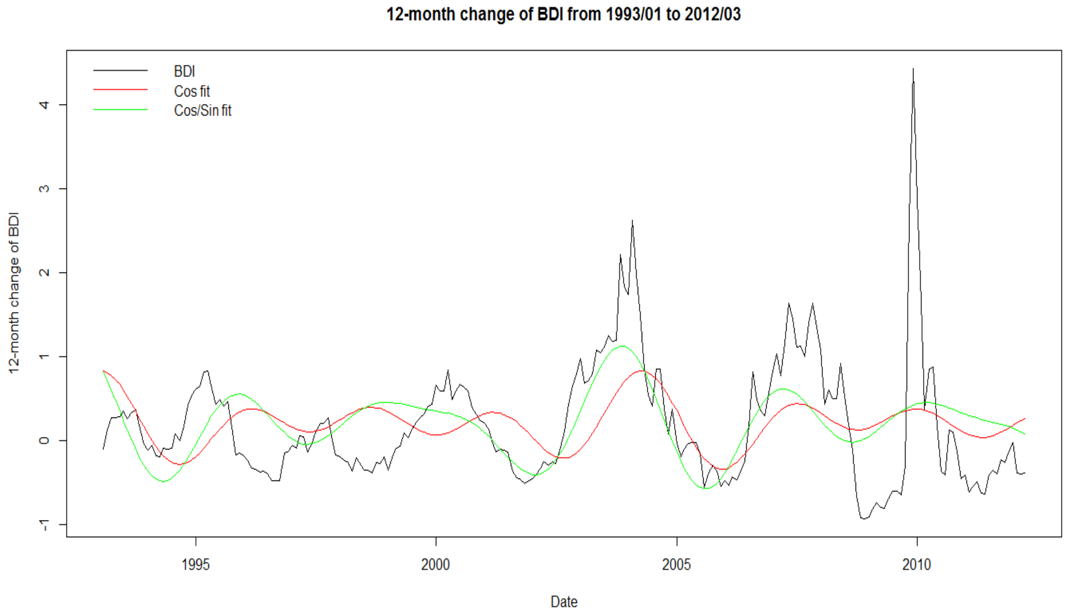

Figure 6: Trigonometric regression of the annual percentage change of the BDI.

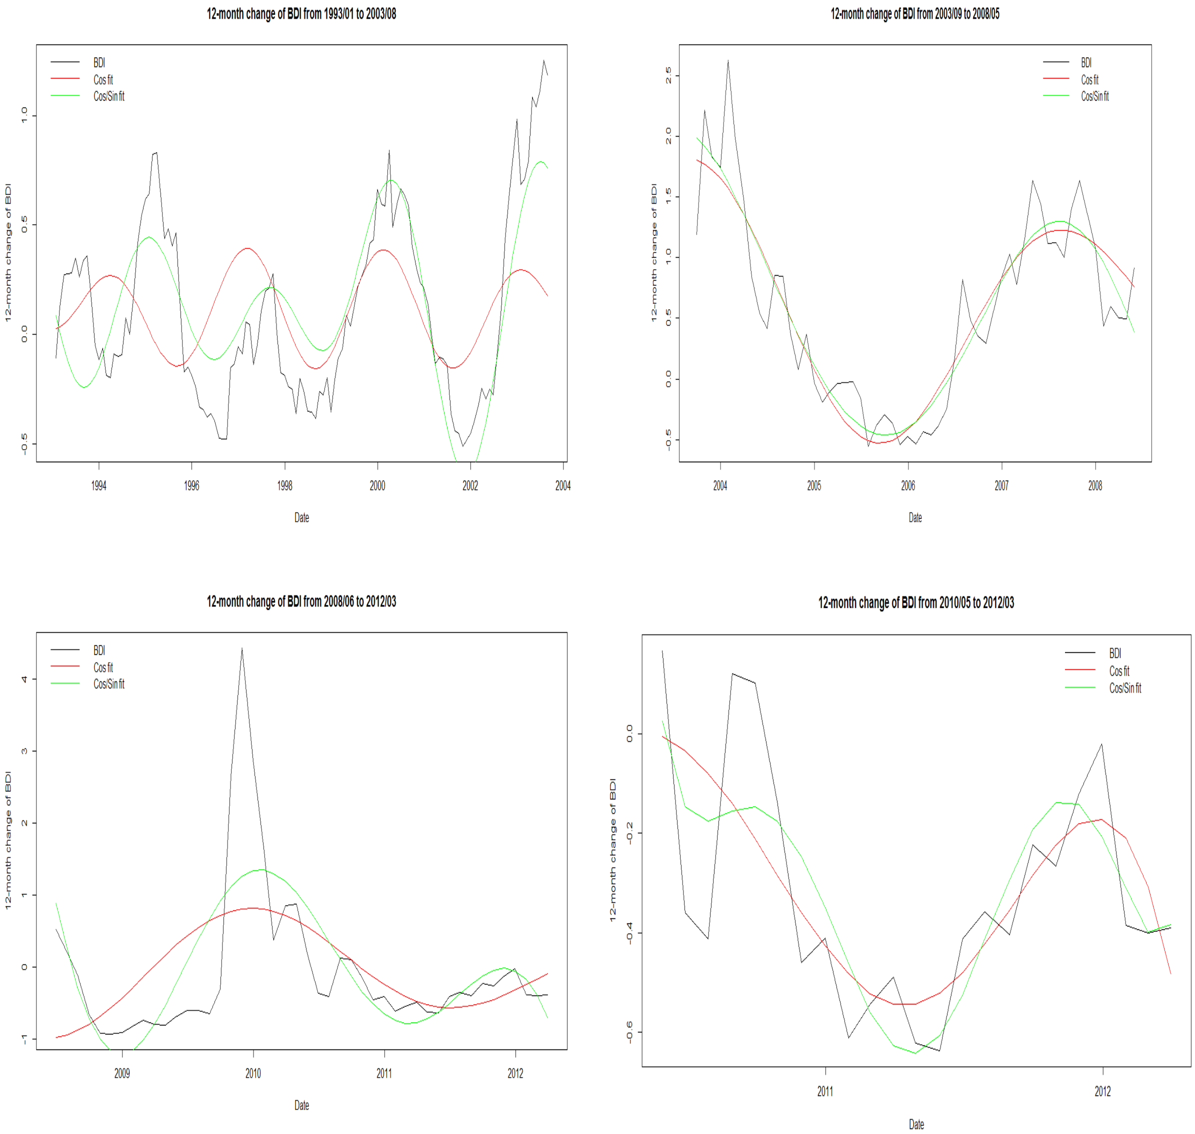

Figure 7: Trigonometric regression of the annual percentage change of the BDI across periods.

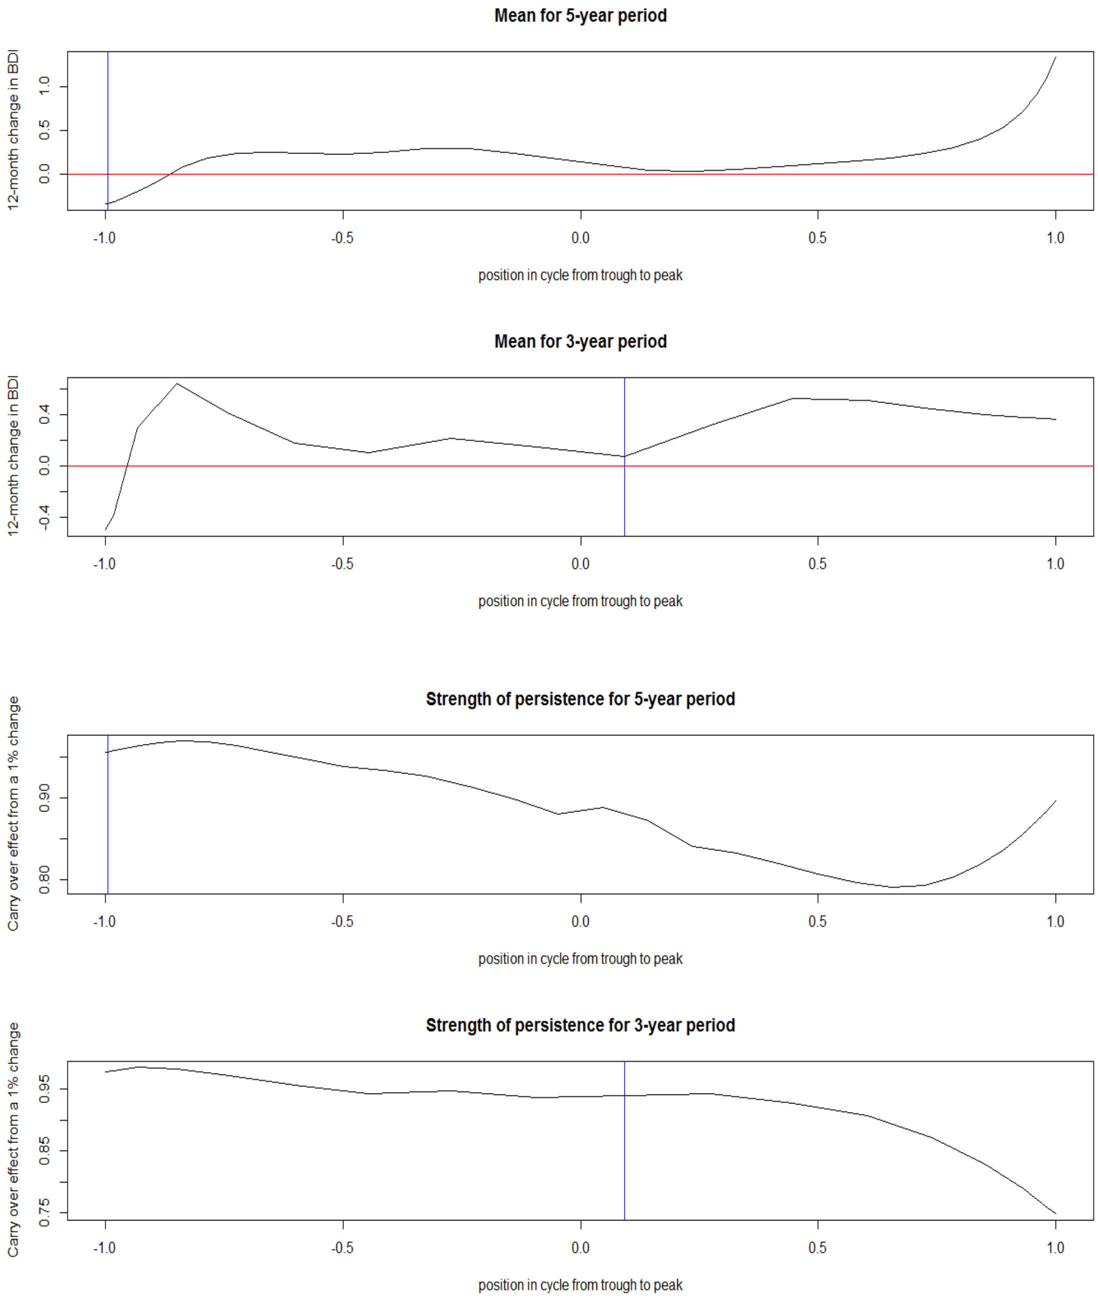

Figure 8: AR and Cycle characteristics of the annual percentage change of the BDI.

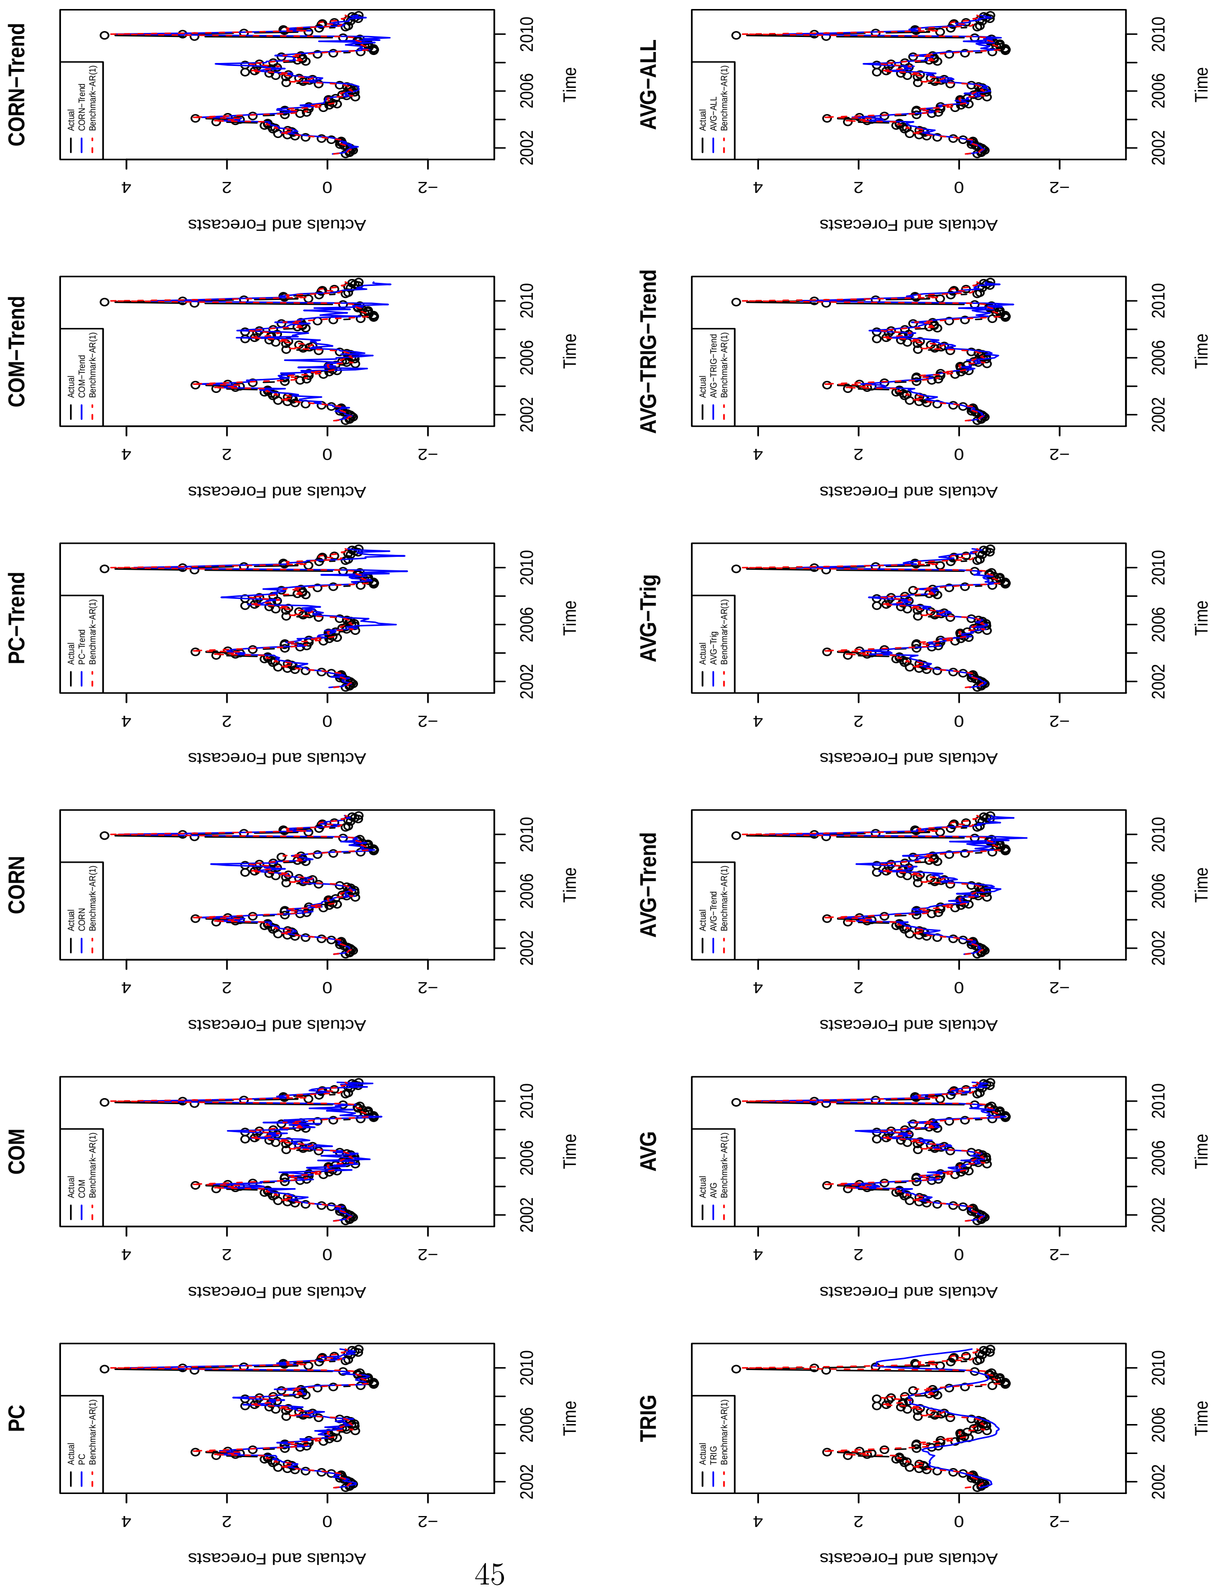

Figure 9: Forecasting evaluation for different models. Rolling window: 90 months, 1-step ahead forecasts.

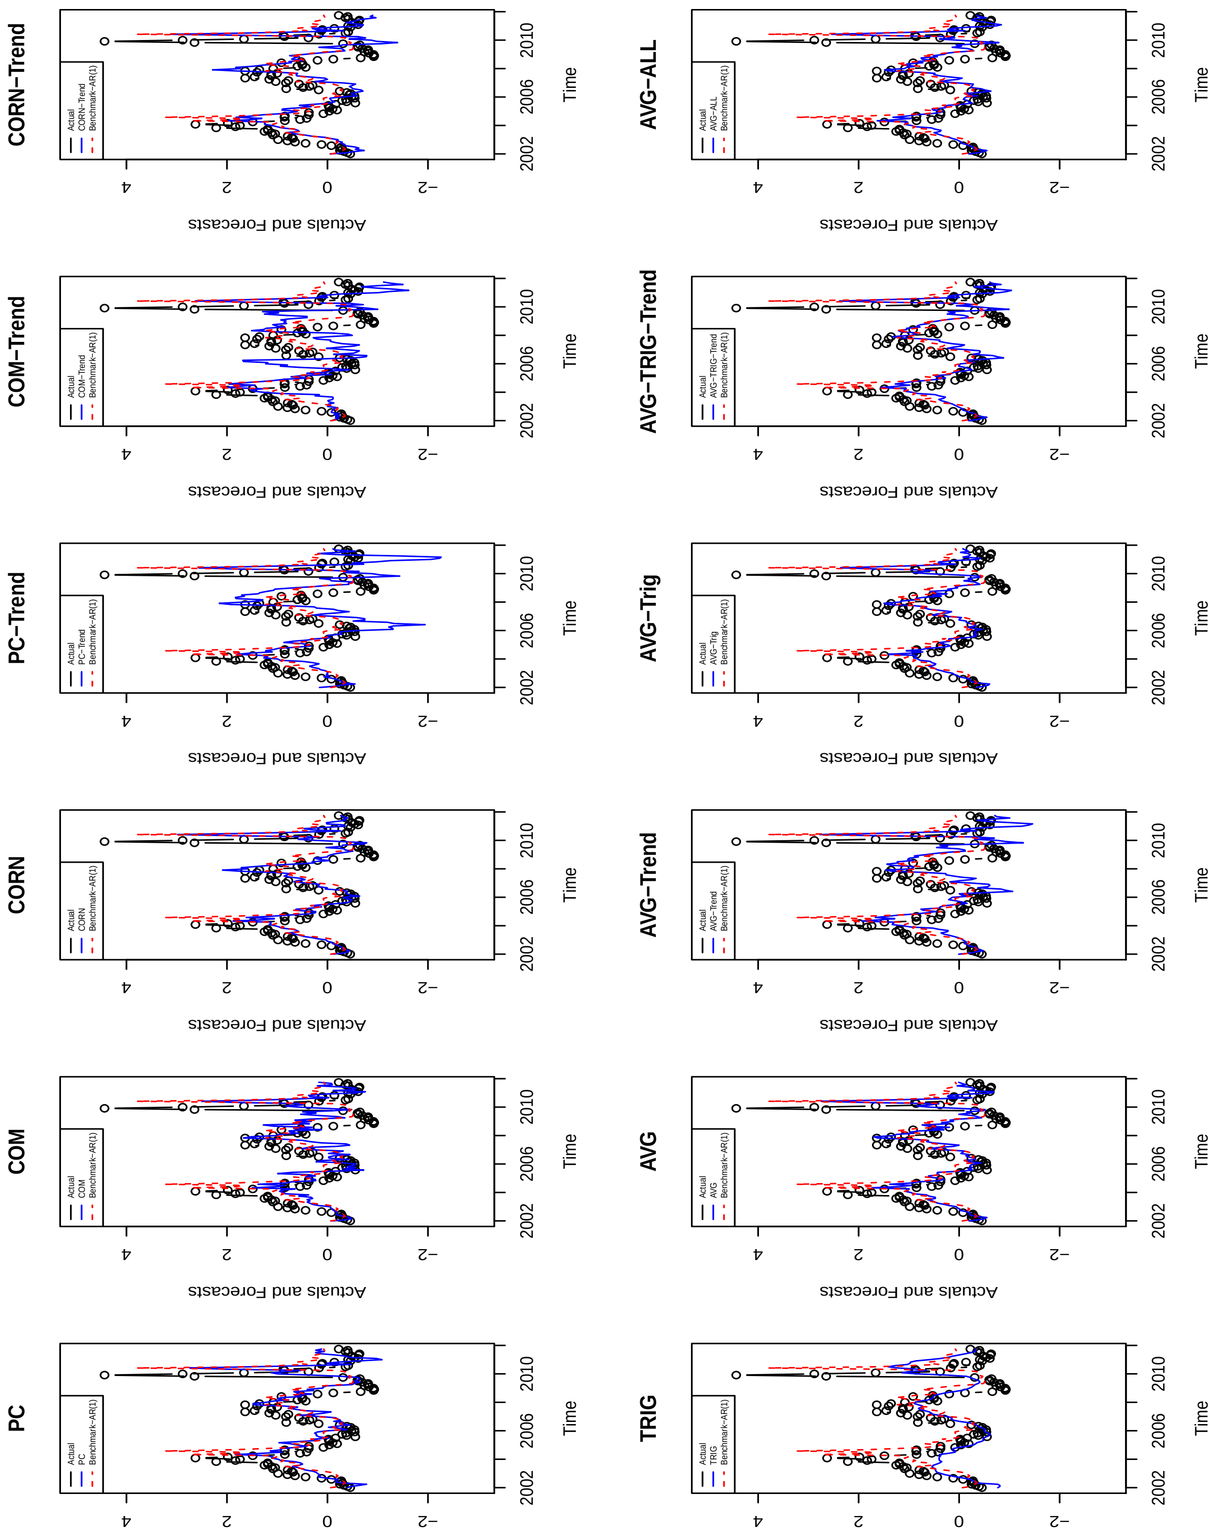

Figure 10: Forecasting evaluation for different models. Rolling window: 90 months, 6-steps ahead forecasts.

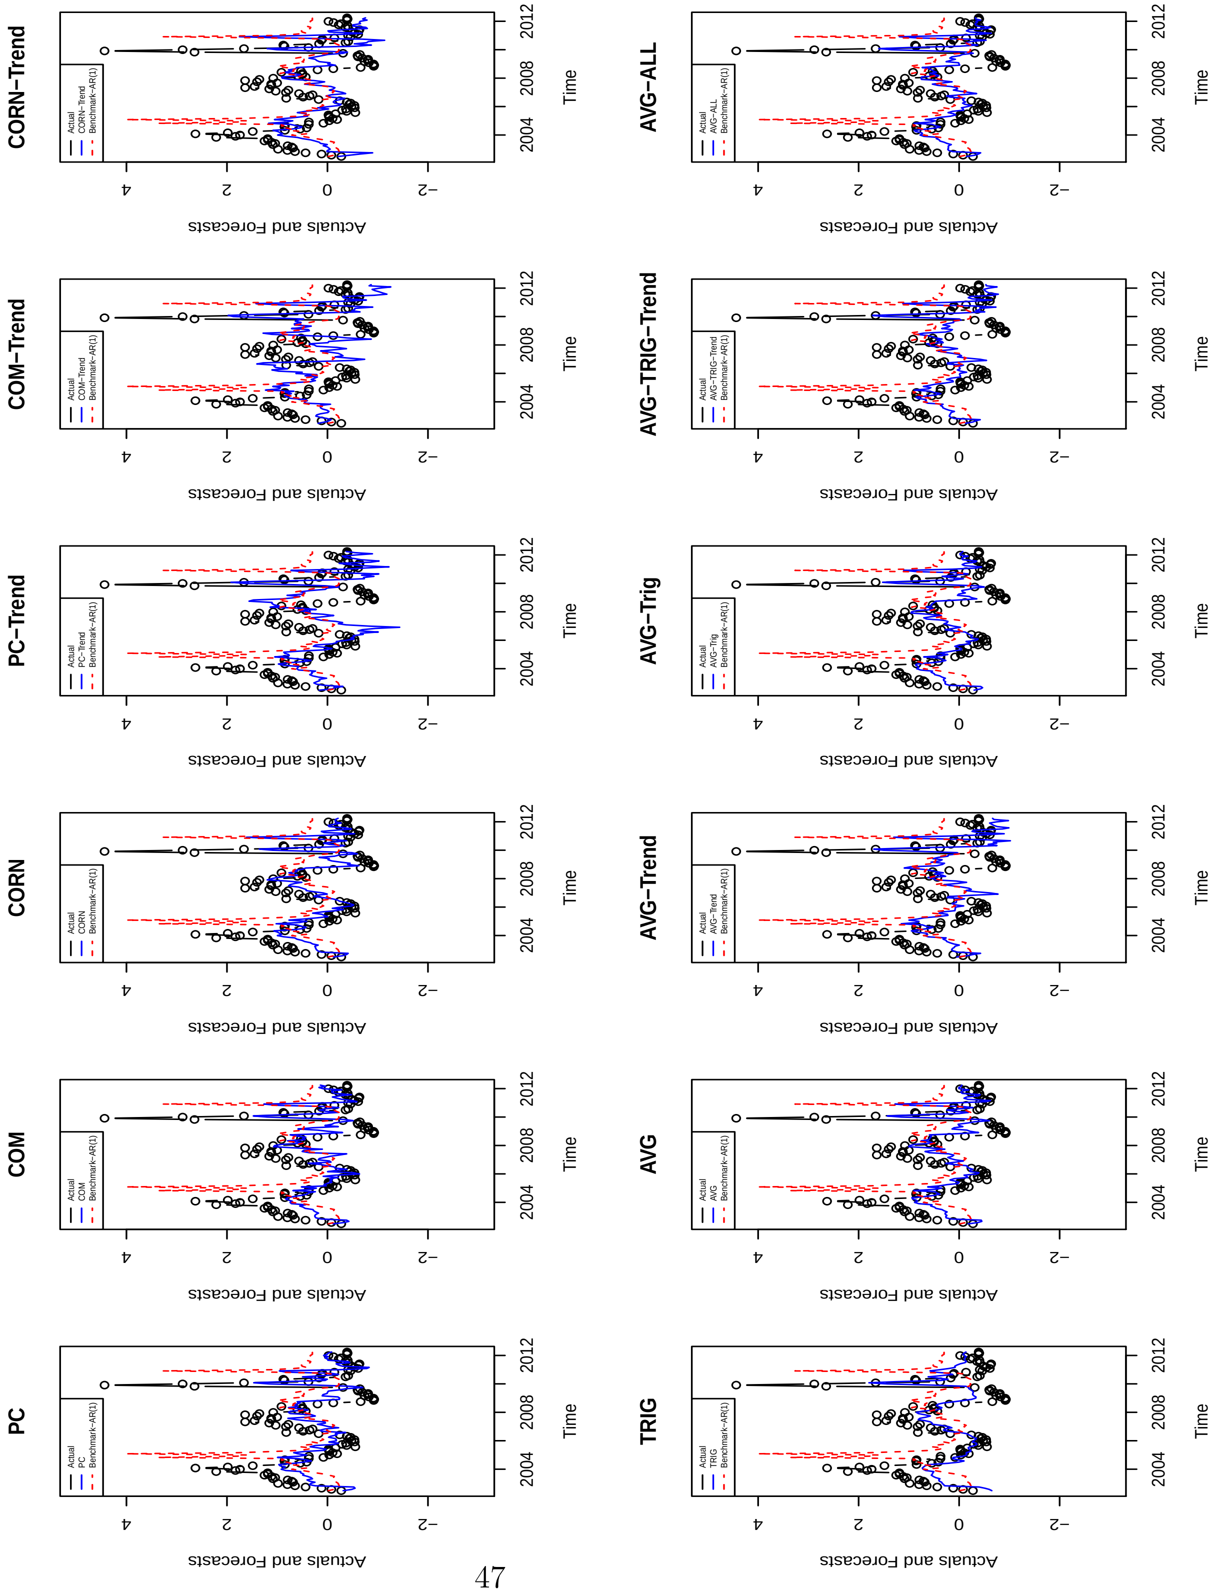

Figure 11: Forecasting evaluation for different models. Rolling window: 90 months, 12-steps ahead forecasts.

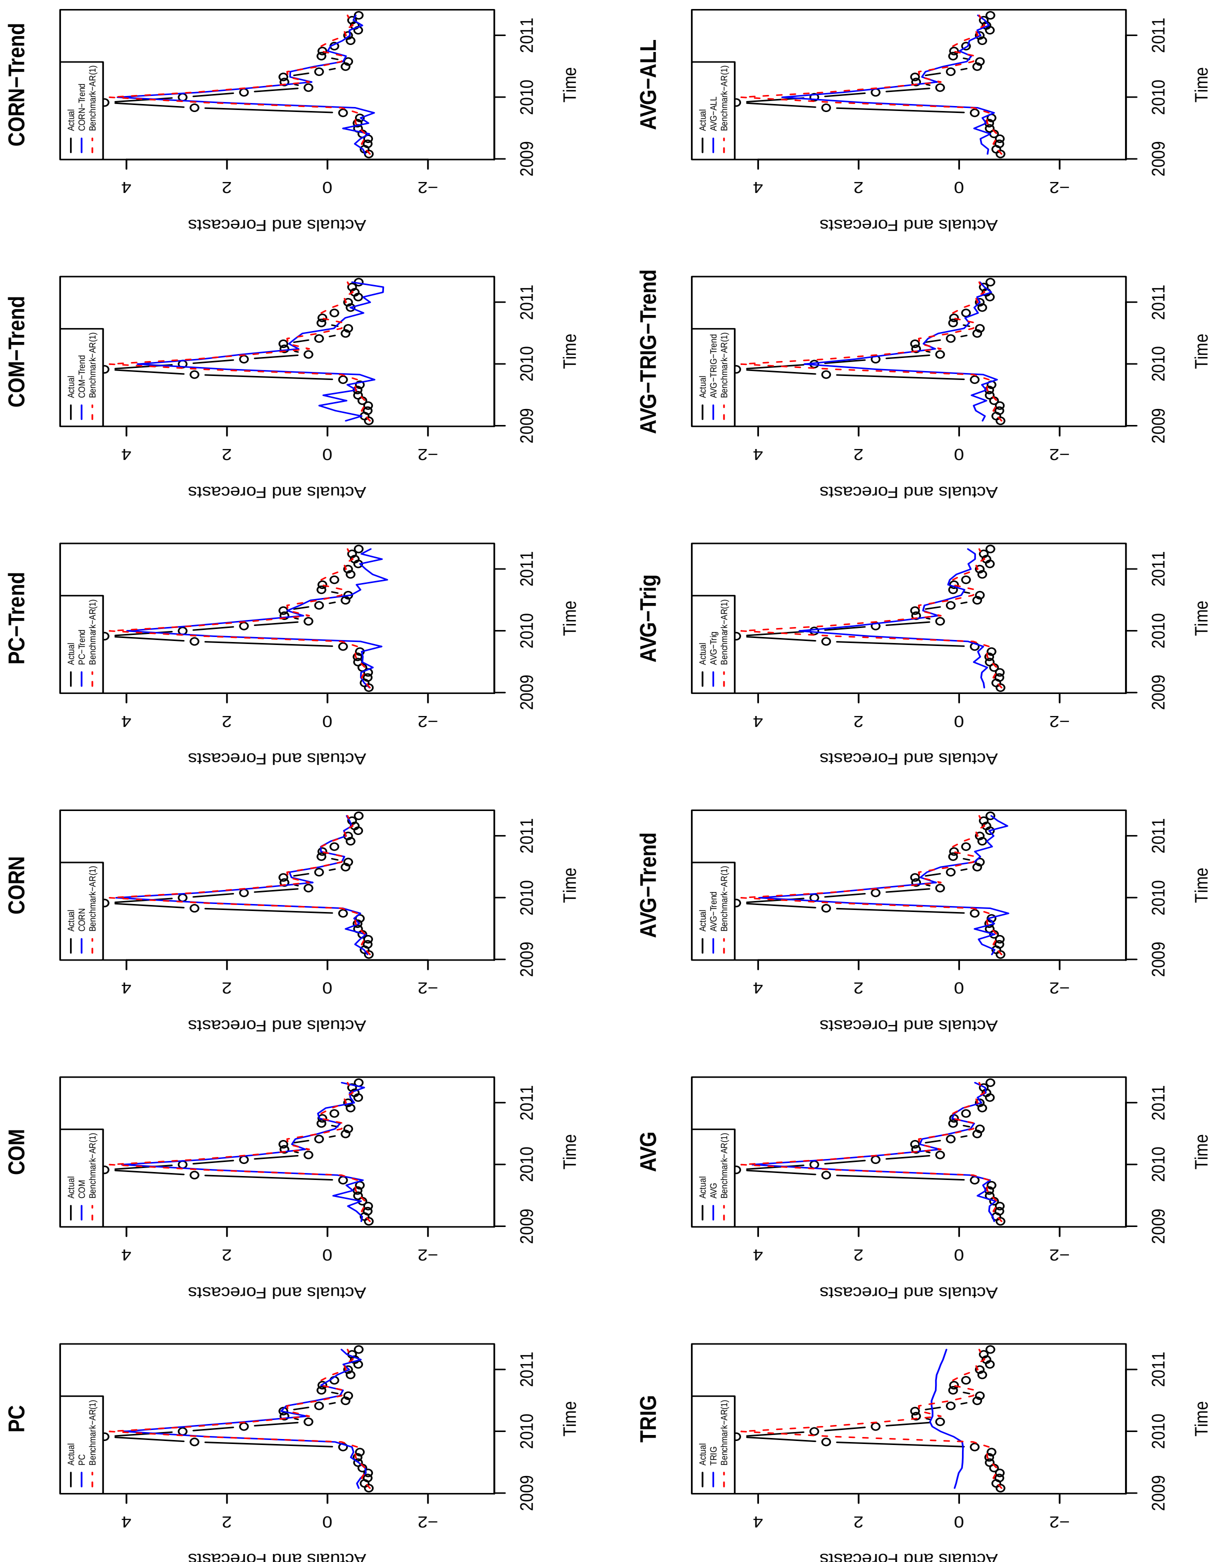

Figure 12: Forecasting evaluation for different models. Rolling window: 180 months, 1-step ahead forecasts.

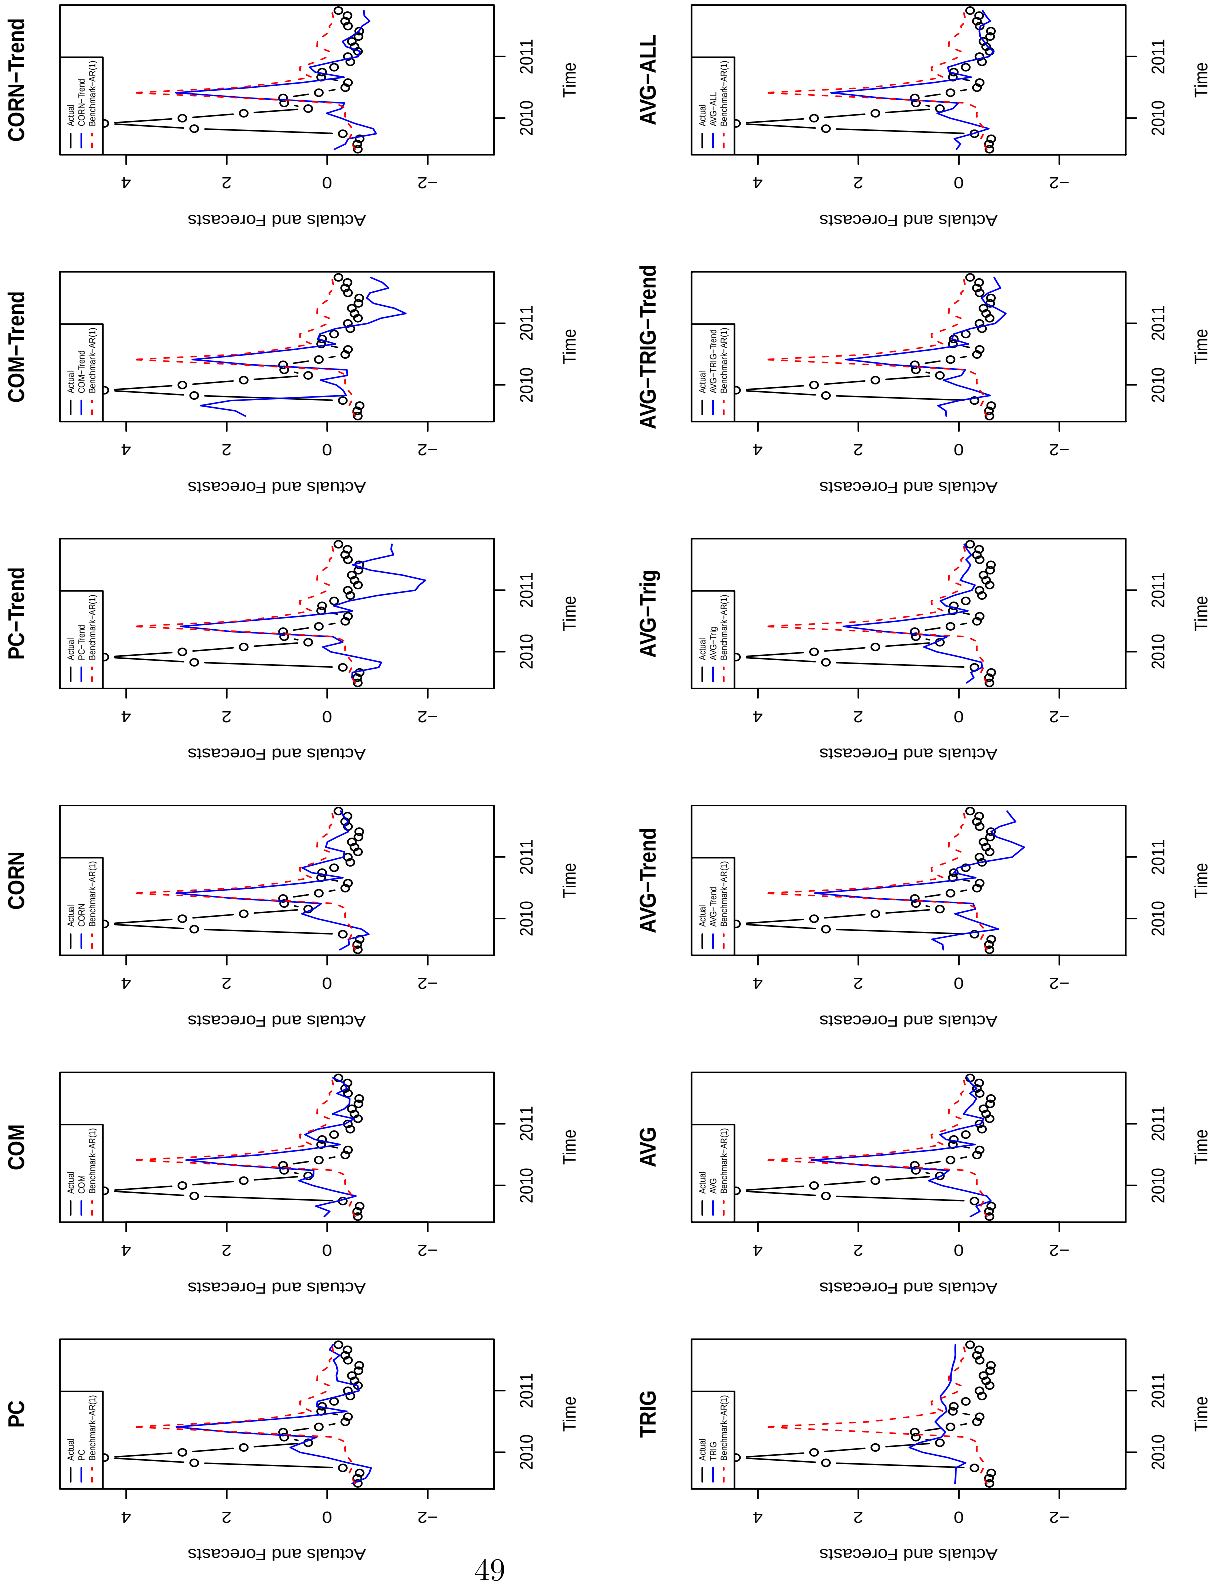

Figure 13: Forecasting evaluation for different models. Rolling window: 180 months, 6-steps ahead forecasts.

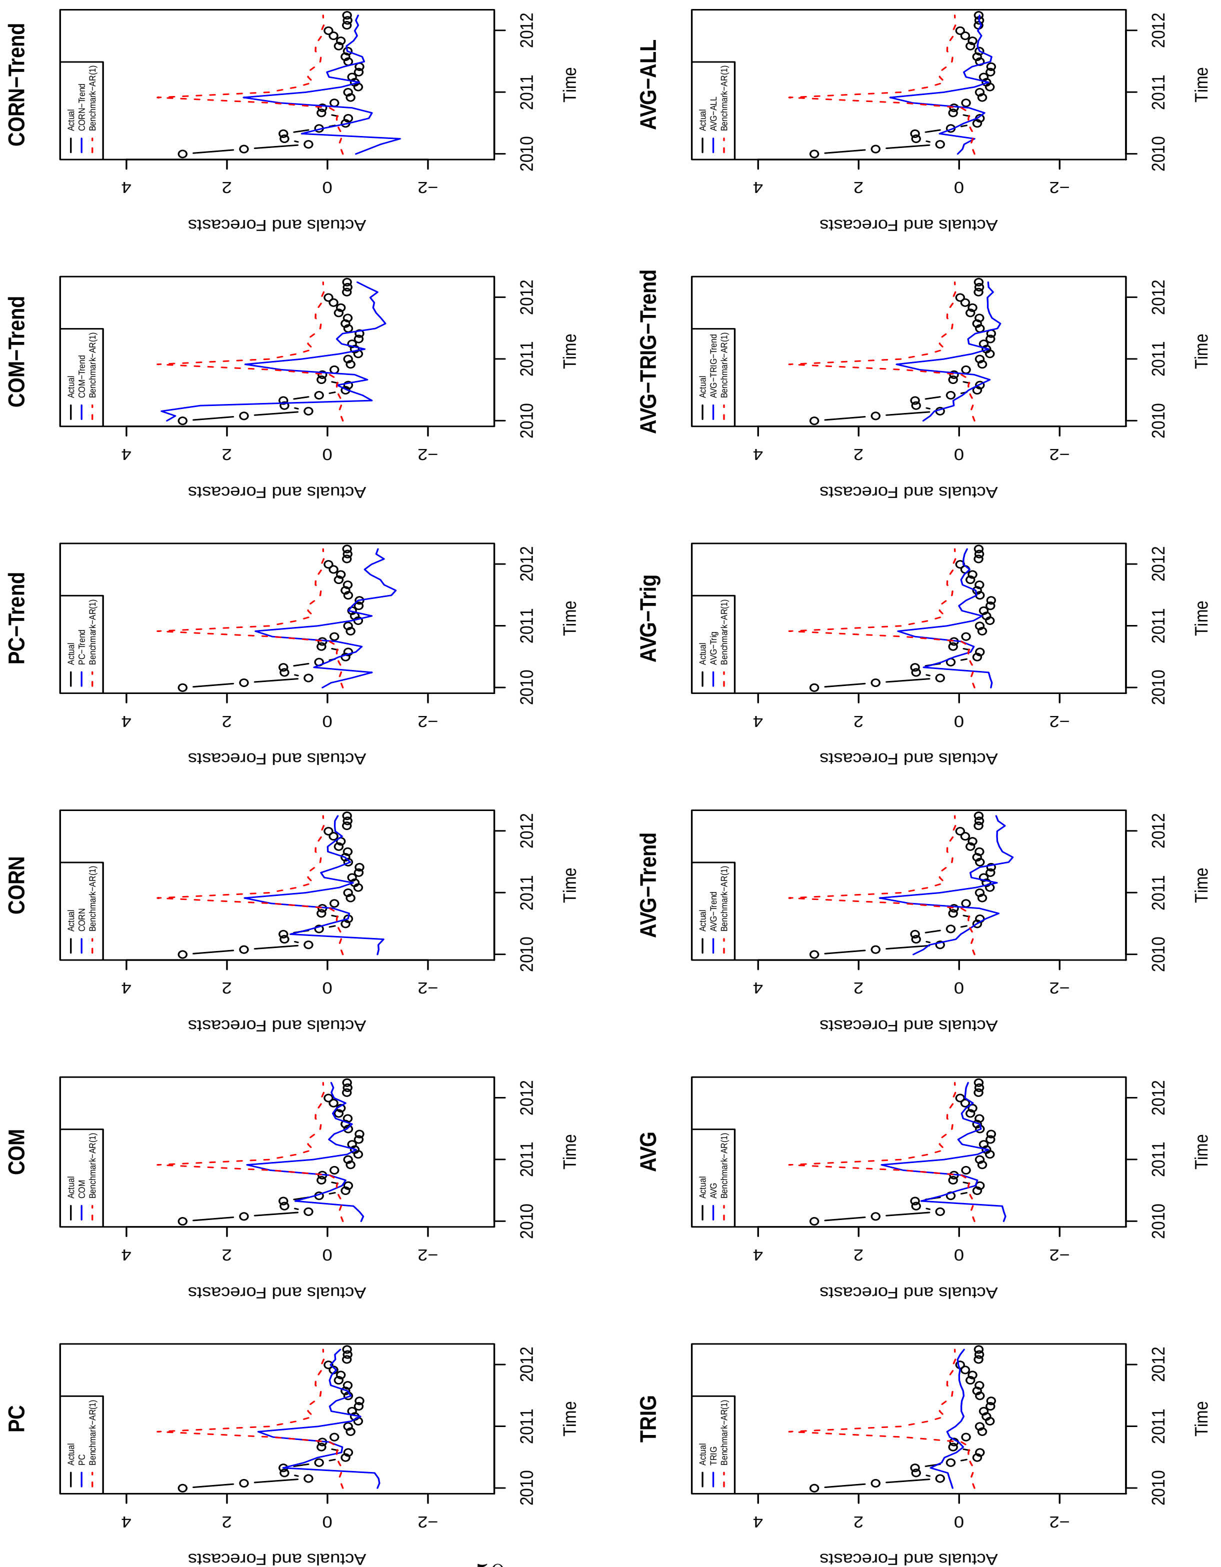

Figure 14: Forecasting evaluation for different models. Rolling window: 180 months, 12-steps ahead forecasts.

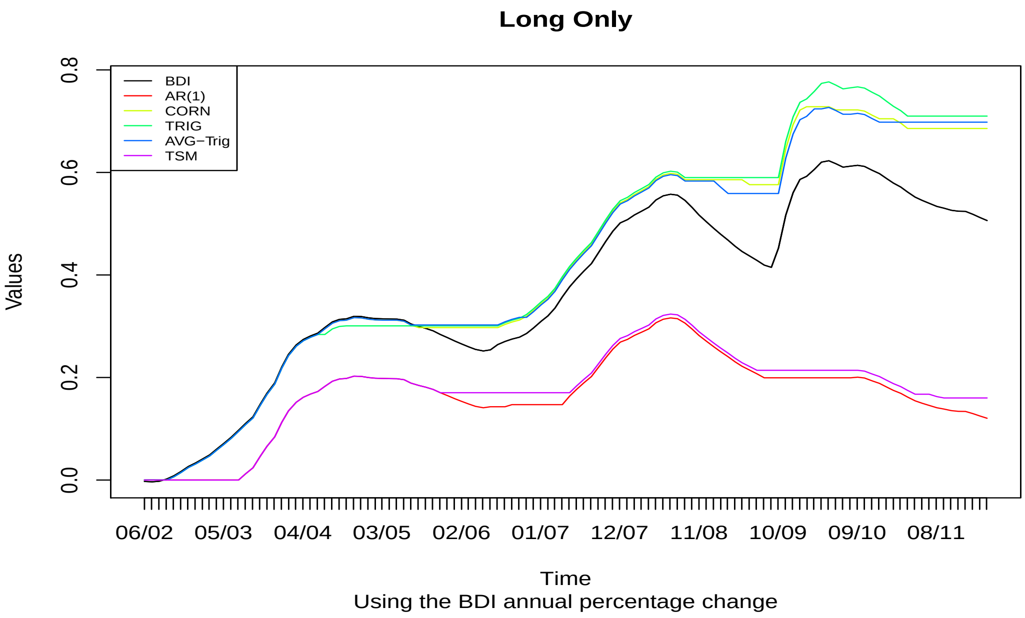

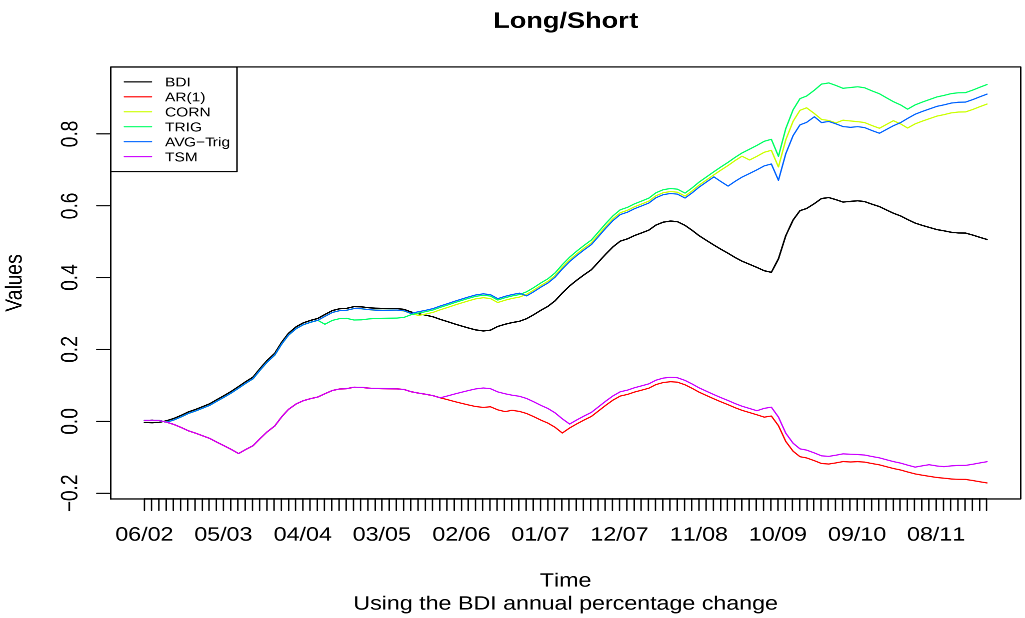

Figure 15: Investing in the BDI. Comparing the cumulative performance of different strategies using 118 evaluation months.

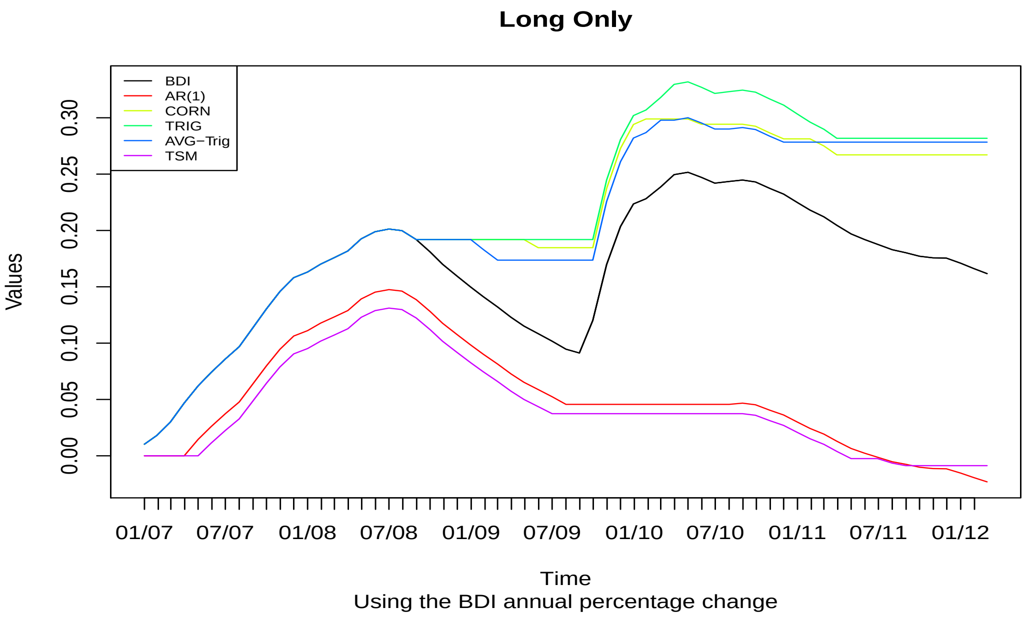

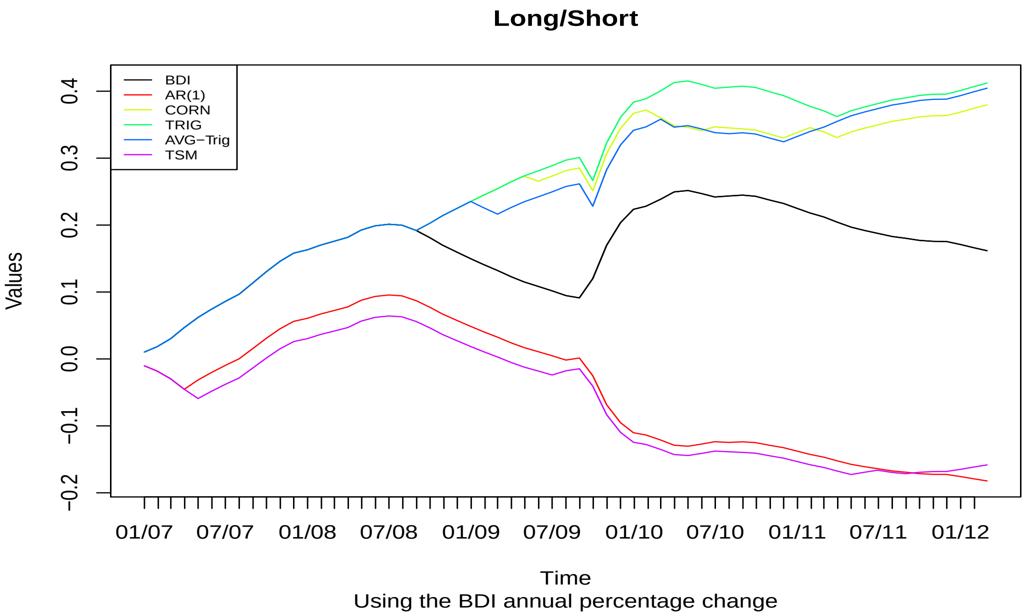

Figure 16: Investing in the BDI. Comparing the cumulative performance of different strategies using 63 evaluation months.