Makridakis Spyros (Orcid ID: 0000-0003-2519-8095)

A Novel Forecasting Model for the Baltic Dry Index Utilizing Optimal Squeezing

Spyros Makridakis[a], Andreas Merikas[b], Anna Merika[c], Mike G.Tsionas[d], and Marwan Izzeldin[d]

Abstract

Marine transport has grown rapidly as the result of globalization and sustainable world growth rates. Shipping market risks and uncertainty have also grown and need to be mitigated with the development of a more reliable procedure to predict changes in freight rates. In this paper, we propose a new forecasting model and apply it to the Baltic Dry Index (BDI). Such a model compresses, in an optimal way, information from the past in order to predict freight rates. To develop the forecasting model, we deploy a basic set of predictors, add lags of the BDI and introduce additional variables, in applying Bayesian Compressed Regression (BCR), with two important innovations. First, we include transition functions in the predictive set to capture both smooth and abrupt changes in the time path of BDI; second, we do not estimate the parameters of the transition functions, but rather embed them in the random search procedure inherent in BCR. This allows all coefficients to evolve in a time-varying manner, while searching for the best predictors within the historical set of data. The new procedures predict the Baltic Dry Index with considerable success.

Key Words: Forecasting; Bayesian methods; Compressed Regression; Baltic Dry Index; Maritime Shipping.

JEL Classifications: C11, C18, C53.

_____
[a] Corresponding Author. University of Nicosia, Institute For the Future (IFF), 46 Makedonitissa Av., Nicosia, 2417 Cyprus
[b] Department of Maritime Studies, University of Piraeus, Karaoli andDimitriou 40, 18534 Piraeus, Greece
[c] Department of Economics, Deree College, The American College of Greece, 6 Gravias St., 15342 Athens, Greece
[d] Lancaster University Management School, Lancaster LA1 4YX, United Kingdom

1. Introduction

The Baltic Dry Index (BDI) has grown into a global economic indicator, a mirror of world trade, whose combined direct and indirect impact on the world economy contributes, through the operation of merchant ships, to about US$380 billion in freight rates, which is equivalent to around 5% of global trade according to UNCTAD (2015). As a weighted average of time charter freight rates, BDI reflects both the supply of cargo ships and the demand for transporting raw and other materials. Until recently, its low levels indicated vessel overcapacity, as well as a slowdown in demand for dry bulk commodities. Since mid 2017 the index has shown definite signs of recovery.

In 2008, the index reached 12,000 points, the highest level in its history; since then, it has entered a spiraling fall, registering a low of 290 in the first quarter of 2016 (Bloomberg, 2016). Yet, the purchase and sale activity of vessels both in 2014 and 2015, according to Clarkson's research, exceeded the 2008 levels, with Europeans being net purchasers and Asians net sellers. This could indicate a turning point, at which economic activity may be entering a new phase of the shipping cycle.

Shipping decisions on the part of ship owners and charterers alike depend on the expected fluctuations of the BDI. Decisions involve entering charter contracts of different durations, switching between spot and time charter operations, improving hedging performance using derivative contracts, as well as whether to invest in newly built or second hand vessels. Since its inception, the BDI has reflected major economic and political events that accounted for its volatile nature. Its time path appears to be close to a random walk, which renders predictions by both academics and practitioners an intriguing challenge. In this paper, we propose a new forecasting model aimed at improving the existing BDI prediction models while reducing uncertainty for decision-makers.

Our study addresses the following research question: given a basic set of predictor variables influencing BDI, how can the dynamics of BDI be used to predict turning points which will take us to a new phase of the shipping cycle, but also reveal short-term upturns or downturns in the market? Valuable contributions in exploring freight rate dynamics and their stochastic seasonal behavior have been made by (Adland, and Cullinane, 2005; Koekebakker, Adland, and Sodal, 2006; Batchelor, Alizadeh, and Visvikis, 2007; and Alizadeh et al., 2007). Further, comparisons of volatility in the dry-cargo ship sector have been conducted by Kavussanos et al., (1999), while the seasonal properties and forecasting in the dry bulk shipping sector have been investigated by (Cullinane 1995; Cullinane, Mason, and Cape, 1999; Kavussanos, and Alizadeh, 2001; and Kavussanos, and Alizadeh, 2002). More recently, (Lin and Wang, 2014; and Kaloupsidi, 2014) explored the nature of fluctuations in world bulk shipping by quantifying the impact of demand uncertainty and time-to-build on

investment and prices. Tsioumas et al., (2017) show that a VARX model is superior in predicting BDI, compared with ARIMA. Overall, most of the major time series methods (Makridakis et al., 1986, Makridakis et al., 1998) have been utilized by scholars in predicting the nature of the BDI, but with limited degrees of success.

Even though substantial work has been done in approximating the time path of freight rates as well as other shipping indicators, when it comes to predicting turning points, what is being referred to as "real" forecasting, practitioners have not been able to rely on the available tools. The BCR model on the other hand, allows a search procedure that narrows down the information set with the aim of better knowing when turning points will occur within the volatile shipping cycle. To do so, the present work draws upon the BCR model by Koop, Korobilis and Pettenuzzo (2016) with two important innovations:

- Apart from the relevant variables and their transformations, we also include transition functions in order to depict smooth or abrupt changes in the BDI. This way we can capture the likelihood of transitions to new states and identify the threshold for "jump effects".
- We do not estimate the parameters of the transition functions, but embed them in the random search procedure inherent in BCR. For this reason, we allow all coefficients to evolve in a time-varying manner and we search for those variables in the predictive set that may generate the most accurate forecasts.

The paper contributes to the forecasting literature in two important ways: first, we show that BDI's complex behavior cannot be captured using time series models by themselves. Instead, we attempt to make predictions indirectly through BCR, in a systematic procedure that can be adapted successfully in the present context. Specifically, the lagged variables account for non-linearity while trigonometric terms account for the cyclical behavior of the BDI by exploiting the properties of Fourier expansions. Furthermore, we use transition functions to account for structural breaks in the series. Second, we compare our model with alternative, popular ones to illustrate its value. These models are AR(1) and ARIMA(p,d,q), ARAR by Carter et al., (2002), ARMAAR by Parzen (1982), DARAR and DARMAAR (differenced ARAR and ARMAAR), and finally the BVAR(m) models by Kadiyala et al., (1997). Our comparative results, which are expressed in the form of MAPE (Mean Absolute Percentage Error), are later presented in Table 4 and confirm the improved performance of the proposed model.

The remainder of this paper is organized as follows: Section 2, presents the proposed model; Section 3, is a brief review of the nature of BDI and of the prior literature; Section 4, outlines the data; Section 5, discusses the main empirical findings and compares them with those of alternative models; Finally, Section 6, presents our conclusions and proposes some suggestions for further research.

2. The proposed model

Suppose y_t is the BDI, the series we need to forecast, and $z_t \in R^{d_z}$ is a vector of covariates thought to be useful for its predictions. Suppose also that $w_t \in R^{d_w}$ is a vector of functional transformations of the basic covariates and their lags, for example squared and higher order terms, interactions, lags of y_t viz. $I_t = \{ y_{t-l}, l = 1, ..., L \}$ and various functional transformations along with interactions with the w_t s -the functional transformations of z_t s, trigonometric terms. Clearly, the number of elements d_w in the final set that includes w_t can be very large. For example if we have a single covariate z_t then we include in w_t terms of the form $z_t^2, z_t^3, ..., z_t^p, ..., z_{t-L}^2, \ z_{t-L}^3, ..., z_{t-L}^p$, terms of the form $y_{t-1}^2, y_{t-1}^3, ..., y_{t-1}^p, ..., \ y_{t-L}^2, y_{t-L}^3, ..., y_{t-L}^p$, terms of the form $\sin(2k\pi z_t), ..., \sin(2k\pi z_{t-L})$, $\sin(2k\pi y_{t-1}), ..., \sin(2k\pi y_{t-L})$, and $\cos(2k\pi z_t), ..., \cos(2k\pi z_{t-L}), \ \cos(2k\pi y_{t-1}), ..., \cos(2k\pi y_{t-L})$ (for $k = 1, ..., K$) and interactions with all previous terms. Moreover, we introduce smooth transition terms of the form:

$$G(c_t; \alpha_c, \delta_{c1}, \delta_{c2}) = \frac{1}{1 + \exp\{ -\alpha_c(c_t - \delta_{c1})(c_t - \delta_{c2}) \}}, \alpha_c > 0, \tag{1}$$

where α_c is a parameter, c_t denotes elements of I_t and z_t and δ_{c1}, δ_{c2} denote parameters of the transition function (Terasvirta, 1998; Tsay, 1998) that allow for possibly non-monotonic behavior. Thus, we can capture smooth transitions to new states, regime switching (depending on the value of α_c) etc. For example, values of α_c close to one indicate slow transition, while values of α_c close to zero indicate fast transitions or (almost) jumps. Moreover, the additional term $(c_t - \delta_{c2})$ is introduced to capture non-monotonic behavior. Further, any variable can be multiplied by the transition function in (1) to introduce possible threshold or "jump" (regime-switching) effects.

Even with few covariates z_t the number d_w of functional transformations that we finally include in w_t, may result into a large *predictive set*. Our basic specification is:

$$y_t = w'_t \beta + u_t, t = 1, \ldots, T, \tag{2}$$

Where we can have, $d_w \gg T$. (d_w = dimensionality of w, T = sample size)

A more flexible model would have been:

$$y_t = w'_t \beta_t + u_t, t = 1, \ldots, T,$$
$$\beta_t = a_o + \beta_{t-1} + \epsilon_t. \tag{3}$$

Where w_t is a vector of predictors and β_t is a vector of coefficients

In this model, we allow all coefficients to evolve in a time-varying manner. Clearly, the model cannot be estimated using standard estimation techniques, Boonstra (2005) so we need to search in w_t for variables or "features", i.e transformations of the existing variables that can generate good forecasts. To this end, we use ideas from Bayes Compressed Regressions (BCR) by Guhaniyogi and Dunson (2015). Thus, we replace (3) with the following:

$$y_t = \tilde{w}'_t \gamma_t + u_t, t = 1, \ldots, T,$$
$$\gamma_t = a + \gamma_{t-1} + \varepsilon_t, \tag{4}$$

Where,

$$\tilde{w}_t = \Phi w_t, \tag{5}$$

Φ is an $d_w \times r$ matrix, γ_t is $r \times 1$ and $r \ll d_w$.

Due to the compression resulting from Φ we call the model "squeezing from the past". Guhaniyogi and Dunson (2015) propose *not* to estimate $\Phi = [\Phi_{ij}]$ but generate random numbers for its elements as follows:

$$P(\Phi_{ij} = \varphi) = \varphi^2,$$
$$P(\Phi_{ij} = 0) = 2\varphi(1 - \varphi),$$
$$P(\Phi_{ij} = -\varphi) = (1 - \varphi)^2, \tag{6}$$

Where $\varphi \in [0.1, 1)$ is a parameter

With a uniform distribution over φ in $[0.1, 1)$ and for reasons of numerical stability, various values of dimensionality reduction r are being deployed in R random searches over the random numbers that comprise the different elements of Φ, subject to the orthonormality constraint $\Phi'\Phi = I$. Guhaniyogi and Dunson (2015) show that BCR has excellent performance in very large-dimensional data sets.

Our final extension to BCR is to replace the dynamic random coefficient specification in (4) with the following:

$$y_t = \tilde{w}'_t \gamma_t + u_t, \, t = 1,...,T,$$
$$\gamma_t = a + \gamma_{t-1} + (I \otimes w_t)\alpha + \varepsilon_t, \tag{7}$$

Where α is an $rs \times 1$ vector of parameters, w_t is a different $s \times 1$ compressed version of w_t, and

$$\tilde{w}_t = \Psi w_t, \tag{8}$$

Where $\Psi = [\Psi_{ij}]$ is another $d_w \times s$ compression matrix for which we have:

$$P\left(\Psi_{ij} = \psi\right) = \psi^2,$$
$$P\left(\Psi_{ij} = 0\right) = 2\psi(1-\psi),$$
$$P\left(\Psi_{ij} = -\psi\right) = (1-\psi)^2, \tag{9}$$

Where $\psi \in [0.1, 1)$ is a parameter

With a uniform distribution over ψ in $[0.1, 1)$, various values of dimensionality reduction r were utilized in R random searches, over the random numbers that comprise the different elements of Φ subject to the orthonormality constraint $\Phi'\Phi = I$. Notice that the dimensionality r and s of Φ and Ψ respectively are allowed to be different.

The purpose of including a set of compressed covariates w_t is to allow better tracking of the time varying coefficients but also better tracking of smooth transitions, regime switches and other changes. For the error term, we assume:

$$u_t \sim N\left(0, \sigma^2\right), \, \varepsilon_t \sim N_r\left(0, \Sigma\right). \tag{10}$$

Our prior has the following form:

$$a \sim N(0, h_a^2 I),$$

$$\alpha \sim N\left(0, h_\alpha^2 I\right). \tag{11}$$

Where I denotes the identity matrix.

The smoothing parameters h_a and h_α can be assigned arbitrarily based on prior beliefs or, according to our practice, can be selected using out-of-sample performance in a hold-out sample consisting of 100 observations.

For σ_u and Σ we use the standard reference or "uninformative" priors (e.g. Zellner, 1971, p. 60 and p. 225).

Our other choices involve the number of powers (p), the number of lags (L), and the number of Fourier terms (K). For the full predictive set w_t, we set $p = 4$, $L = 4$ and $K = 5$.

We arrived at these values starting from a more general configuration setting p=8, L=12 and K=10 as an upper bound. We then used the marginal likelihood function (see equation (12) below) to determine optimal values. Our choice lead to overfitting, i.e. we do not actually need so many variables. We repeated the same exercise until, on the basis of the marginal likelihood criterion (standard use in Bayesian analysis), we finally ended up with p=L=4 and K=5. These values maximize the marginal likelihood.

We allow for full interactions of all (basic) variables in w_t (see equation (5) above). Parameters α_c, δ_{c1} and δ_{c2} in the transition function (1) are drawn randomly in the context of R searches. All are drawn from standard uniform distributions, but series c_t in (1) is normalized to lie in $[0, 1]$.

Given that parameters ψ and φ are drawn from uniform distributions in $(0.1, 1]$ we perform $R = 10^6$ searches over the different random elements of Φ and Ψ. We select the model that performs best in terms of out-of-sample behavior using a hold-out sample of 100 observations.

Bayesian inference can be implemented using a standard Gibbs sampler in (4). We use 15,000 passes, the first 5,000 of which are discarded to mitigate the possible impact of start-up effects. Since we do not estimate Φ or Ψ the model can be executed easily, even with a large number of core processors. To maximize the marginal likelihood and select a model during the training period, we use the Laplace approximation by Lewis and Raftery (1997) that can be implemented in a straight manner, since the marginal likelihood is given by:

$$M(X) = \frac{L(\theta; X) p(\theta)}{p(\theta \mid X)}, \tag{12}$$

For any model whose parameter vector is $\theta \in \Theta \subseteq R^d$, the prior is $p(\theta)$, the posterior is $p(\theta \mid X)$ and the likelihood is $L(\theta; X)$ from data X. To estimate the denominator, which is unknown, we use a multivariate normal approximation. If the posterior mean and covariance matrix are $\overline{\theta} = E(\theta \mid X)$ and $\overline{\Omega} = E\left[(\theta - \overline{\theta})(\theta - \overline{\theta})' \mid X\right]$, it is not difficult to establish that an approximation to (12) is

$$\log M(X) \simeq \log L(\bar{\theta}; X) + \log p(\bar{\theta}) + \frac{d}{2} \log |\bar{\Omega}|. \tag{13}$$

The required posterior moments are estimated using MCMC (Markov Chains Monte Carlo) methods.

3. The Nature of BDI and Prior Literature

The origin of BDI goes back to 1744, at a London coffeehouse named "Virginia and Baltik," then known as a meeting place where merchants and ship owners negotiated agreements for shipping goods and products (Harlaftis et al., 2012). The modern BDI was launched in 1985 and in its current form covers freight rates on 26 shipping routes and four vessel sizes. It focuses solely on the shipping of dry-bulk goods; among its determining factors are the vessel size and the carried cargo (Silverberg et al., 2009).

Even though, it is generally accepted that what determines BDI is the law of supply-and-demand, it is also true that BDI is influenced by several other factors (Stopford, 2009), such as commodity demand, seasonal variations, bunker prices, and choke points (ICS, 2016). The existing number of vessels and those newly delivered determines the vessel supply, minus the number of cargo ships that are withdrawn. Commodity demand, on the other hand, is influenced by that of raw and other material used in the production of intermediate and final goods, as well as the demand for grains. Further, major economic and political events affect demand and influence the time path of BDI, while seasonal variations and weather can have a large impact on the shipping markets. For instance, unfavorable weather conditions may potentially affect the entrance to ports (e.g., ice in the ports), may cause reduced river levels, or may affect the size of harvest that correlates to the demand for raw material. In addition, bunker prices account for "between one quarter and one third of the cost of running a vessel" (Baltic Exchange, 2016), while choke points (the relatively narrow shipping lanes of the Suez and Panama canals) can become an important determinant of BDI volatility, as they may delay the smooth passage of ships, thus changing supplying patterns (Baltic Exchange, 2016). Moreover, piracy is another determining factor on world shipping, with an estimated annual cost that ranges from seven to twelve billion US dollars (Bowden et al., 2010).

The literature on freight rates is rich. Several papers have addressed various aspects of the behavior and the time series properties of such rates. Driehuis (1970) first investigated and modelled liner freight rates. Beenstock and Vergottis (1989a, 1989b) built an econometric model for the world tanker market and the dry bulk market, while Randers and Goluke (2008) claimed that only time

patterns of oil tankers freight rates may be forecasted. The model they used combined both transport demand and consumption in the oil market; it provided reasonable forecasts for a period of one to four years, with an uncertainty of plus/minus six months.

The volatility in the dry-cargo ship sector has been studied by Kavussanos (1996), while its seasonal properties and forecasting were investigated by (Cullinane, 1995; Cullinane, Mason and Cape, 1999; and Kavussanos and Alizadeh, 2001). Further, there were more studies by (Adland and Cullinane, 2005; Koekebakker, Adland and Sodal, 2006), while Batchelor, Alizadeh and Visvikis (2007) considered the properties and the dynamics of the BDI since 1990. They showed that spot freight rates are forecastable and Forward Freight Agreements (FFA's) help to predict spot freight rates. They found that forward rates converge strongly on spot rates, and that Vector Error Correction models (VECM) provided the best in sample fit. In addition, they asserted that with some degree of speculative efficiency, forward rates help in the forecasting of spot rates using ARIMA or VAR models for the prediction of forward rates. The relationship of the BDI with trade and income was investigated by Lin and Wang (2014), and by Kaloupsidi (2014), who explored the nature of fluctuations in world bulk shipping by quantifying the impact of demand uncertainty and time-to-build on BDI prices. More recently, Papailias et al., (2016) showed that the BDI has a cyclical pattern that has been stable, except for a period after the 2008 crisis, which suggests that such a pattern has implications for improved forecasting.

To illustrate the volatility of BDI, we have traced its historic values since 1947[1]. Its time path indicates that it is subject to major political and economic events that affect seaborne trade (see Figure 1 showing its value between 1947 and 2016).

[please insert Figure 1 here**]**

It appears from the graph above that the shipping industry has been historically affected by the major political and economic events of the world. The closure of the Suez Canal in 1956-57 and 1967, for instance, forced ships to go around the Cape of Good Hope, which added thousands of miles to voyages and dramatically increased fuel costs, thus in turn causing a sharp increase in freight rate prices. Likewise, in the early sixties, the increased demand for grains from China and the USSR led the BDI onto a higher level that was maintained amidst various wars, the Asia crisis, as well as the Dot.com collapse.

---

[1] We have been allowed access to the records of a globally leading shipping company since the beginning of the 20th century.

It was intense demand, prior to the crisis, that led to an oversupply of ships afterwards. More specifically, China's demand for raw material was so high that it provided the main impetus for the introduction of a large number of new ships to the shipping market (German, et. al., 2012). As it usually takes two to three years between the placing of an order and the vessel being put into service, a large number of new ships arrived in the market during and after the 2008-9 Great Recession. For instance, in 2010 the global fleet increased by 23%, creating an over-supply of cargo ships that combined with the declining demand, contributed to the fall of BDI to one of its lowest levels (Oomen, 2012).

Following the financial crisis, two developments occurred. One was the collapse in demand, as world exports stopped growing (World Bank, 2016; OECD, 2016), and the other was the structural shift of the Chinese economy away from heavy manufacturing, which resulted in the drop of its economic growth rate from 20% to close to 7%. These changes affected the Baltic Dry Index that fell to a 30-year low of 290 in February of 2016. There was a brief rebound in the summer of the same year, but the index continued its decline afterwards, reflecting vessel oversupply and lower growth in the demand for commodities. Since mid-2017, it has started to regain some of its lost ground.

4. Variables used and descriptive statistics

The data employed in our study (see Table 1 below) covers the period from January 1990 to September 2016. It was collected mainly from Clarksons, a web-based shipping database. The data on the Baltic Dry index was gathered from the records of a leading shipping company that covered the annual period between 1947 and 1984; data from the period 1985 to 2016 was also derived from Clarksons. For consistency reasons, our study covers the period from 1990 to September 2016, in a monthly frequency. Panel A of Table 1 provides the definition of the variables used: We collected the main sub-indices of the dry bulk sector which reflect demand and supply conditions by segment and vessel type. The Baltic Capesize Index (BCI) reflects conditions in the Capesize segment, which consists of vessels between 80,000-100,000 DWT. The Baltic Panamax Index (BPI) reflects conditions in the Panamax segment, which consists of vessels between 65,000-80,000 DWT. The Baltic Handysize Index (HSI) reflects conditions in the Handysize segment, which consists of vessels between 10,000-40,000 DWT. We have also included the fleet size - as the main supply factor, in vessel numbers as well as in DWT - of the smaller segments of the dry bulk sector, i.e. Handysize and Handymax (40,000-65,000 DWT). The inclusion of IRONORE imports represents demand conditions

(demand for shipping services is a derived demand) in the dry bulk sector. Lastly, as a proxy to the phase of the shipping cycle we have included the second hand vessel price (SHP).

Panel B in Table 1 shows the main descriptive statistics of the variables included in the model. We observe for all indices, BDI, BCI, BPI, HSI, but also for the second hand vessel price, SHP, and IRONORE imports large dispersion from their mean values. This shows the highly volatile nature of the market over the period we investigate.

Panel C in T able 1, shows correlations between the variables included in the model. It appears that correlation over the entire period between the indices of the different types of vessels is over 95%, implying that there is not much scope for diversification within the dry bulk sector. Some other correlations, however, are smaller.

5. Empirical results

We use monthly observations from 1990 to 2016. The model is estimated from January 1990 to December 2006. Our measures of predictive accuracy are compared to an AR(1) benchmark. To examine the precision of the h-step-ahead point forecasts for a given model i, we use the ratio of mean-squared forecast errors by Koop, Korobilis and Pettenuzzo (2016):

$$MSFE_{i,h} = \frac{\sum_{\tau=t_1}^{t_2-h} \hat{e}_{i,\tau+h}^2}{\sum_{\tau=t_1}^{t_2} \hat{e}_{AR1,\tau+h}^2}, \qquad (14)$$

Where t_1 and t_2 denote the start and end of the out-of-sample period, and $\hat{e}_{i,\tau+h}^2$, $\hat{e}_{AR1,\tau+h}^2$ are the squared forecast errors at time τ from models i and AR(1) respectively, when the forecast horizon is h.

In Table 2 we compare the out-of-sample forecasting performance of several alternatives relative to the benchmark, the simple AR(1) model. The alternatives differ in the contents of the predictive set as shown in Table 2. The basic set includes lags of BCI and BPI, HAMDWT, HAMNO, HANDNO, HANDWT, HSI, SHP, and IRONORE.

The results (MSFEs) support the idea that squeezing from E (which includes lags of basic sets, squares and interactions, plus trigonometric terms ($K = 5$) and transition functions) performs very well contrary to sets A, B, C, D which do not perform well relative to the benchmark AR(1) model.

Following Geweke and Amisano (2010), we also use the log predictive likelihood differential between model i and the benchmark AR(1) model:

$$ALPL_{i,h} = \frac{1}{t_2 - t_1 - h + 1} \sum_{\tau=t_1}^{t_2-h} \left(LPL_{i,\tau+h} - LPL_{AR1,\tau+j} \right),$$ (15)

Where $LPL_{i,\tau+h}, LPL_{AR1,\tau+j}$ are log predictive scores for model i and benchmark AR(1) respectively at time $\tau + h$, i.e. the log of the h-step-ahead predictive density evaluated at the outcome. Table 3 below shows that positive values of ALPL indicate that on average model i produces better forecasts compared to the benchmark.

These results show that squeezing produces substantial improvements in the log predictive likelihood differentials relative to the benchmark AR (1). Both Tables 2 and 3 demonstrate that the forecasting performance depends a lot on including transition functions, rather than trigonometric terms or past values of the basic set.

The actual forecasting performance of model E is presented in Figure 2 for the period 2007:1 to 2016:6 and in Figure 3 for the period 2014:1 to 2016:6, showing detail.

5.1 Sensitivity analysis for the hyperparameters

As ψ is selected optimally, the issue remains of whether our results are sensitive to the hyperparameters h_a and h_α which have been selected based on a holdout sample. We change these hyperparameters by λh_a and λh_α where λ is uniformly distributed in the interval (0.1, 10). We select 10,000 different λ and we repeat our forecasting exercise. We proceed by performing sensitivity tests on the results reported on Tables 2 and 3, with Table 4 and Table 5 reporting maximum absolute percentage changes. The results indicate that the maximum absolute percentage changes are trivial, meaning that changing the hyperparameters widely relative to their optimal values (derived from the hold-out sample) does not produce important differences in the final results.

Figures 2 and 3 show the forecasting performance of model E within and out of sample respectively. It appears that major turning points are being captured.

In Figure 4 we report the log predictive likelihood of Model E (this time not relative to the benchmark AR(1) model). The log predictive likelihood is normalized to have the same mean with log BDI for visual clarity. This Figure shows that the log predictive likelihood increases in the middle of 2008, before the sub-prime crisis, and continues to increase almost monotonically afterwards, suggesting that the model performs rather well and captures significant aspects of the data. Figure 4 below presents the value of log predictive likelihood. It shows how the log predictive likelihood behaves over time relative to the data (It is only shown for completeness).

5.2 Comparison of BCR with Alternative Models

In Table 6, we present a comparison between our model (normalized Mean Absolute Predictive Error) and a number of alternative ones accepted in the forecasting literature. The comparison is based on sliding forecasts of 36 months, starting from the end of 2016. The orders p,d,q in ARIMA(p,d,q) are selected using the Schwarz BIC criterion. The ARAR model, Carter et .al., (2002), is $y_t = \alpha_t + \beta_t y_{t-1} + u_t$ where α_t and β_t follow a joint VAR scheme. The ARMAAR model, Parzen, (1982), is $y_t = \alpha_t + \beta_t y_{t-1} + u_t + \gamma_t u_{t-1}$ where α_t , β_t and γ_t follow a joint VAR scheme. DARAR and DARMAAR are as ARAR and ARMAAR but formulated in first differences. BVAR(m) is a Bayesian VAR model consisting of m variables, Kadiyala et al. (1997). The six variables we consider for BVAR are BDI, BCI, BPI, SHP, NB, HSI and they enter in BVAR(2), BVAR(4), BVAR(6) in that order. The BVAR uses a Minnesota prior. TVP-VAR is a time-varying parameters VAR, where parameters follow a random walk, see Koop and Korobilis (2013). For all Bayesian models, we use a Gibbs sampler with 15,000 passes, the first 5,000 of which are discarded to mitigate possible start-up effects.

Finally, we present a more elaborate model that searches through all nonlinear combinations of the predictive set, rather than the "simplified" sets A-E in Tables 3 and 4. The model is presented in Table 7.

In Table 7, lags, lags^2 and lags^3 denotes powers of lags of the predictors (reported in the first column). For example, we need 1, 2 and 3 lags of BDI, and 1 lag of BDI squared. From BPI, we need lags 2, 3, 4, squared lags at periods 1 and 2 and the first lag of BPI to the third power. We also need interactions between lagged variables. Moreover, we need $DI_{t-1}BDI_{t-2}$ which we denote by (1, 2) in the fifth column and $BDI_{t-1}BDI_{t-3}$ which we denote by (1,3). Moreover, we need $HAMNO_{t-2}HAMNO_{t-2}$ denoted by (2, 3), $HAMDWT_{t-1}HAMDWT_{t-2}$ etc.

Additionally, we need trigonometric terms of the form sin(kπz) and cos(kπz), where the values of k (a parameter) and z (a variable) are to be determined. For example, in terms of BDI we need $\sin(\pi BDI_{t-1})$, $\sin(\pi BDI_{t-1})$, $\sin(2\pi BDI_{t-2})$ and $\cos(\pi BDI_{t-1})$. Trigonometric terms are needed from BDI, BCI and BPI, which are the basic indices in our analysis. Their relevance means that cyclical features are captured which are important in short-term prediction.

Finally, we need transition functions as in (1). From BCI, for example, we need terms of the form

$$\frac{1}{1+\exp\{-0.80(BCI_{t-1}-\delta_{c1})(BCI_{t-1}-\delta_{c2})\}} \quad \text{and} \quad \frac{1}{1+\exp\{-0.80(BCI_{t-2}-\delta_{c3})(BCI_{t-2}-\delta_{c4})\}}. \;).$$

From IRONORE, we need terms of the form

$$\frac{1}{1+\exp\{-0.90(IRONORE_{t-1}-\delta_{c5})(IRONORE_{t-1}-\delta_{c6})\}} \quad \text{and}$$

$$\frac{1}{1+\exp\{-0.90(IRONORE_{t-2}-\delta_{c7})(IRONORE_{t-2}-\delta_{c8})\}}.$$

In terms of economic significance, the use of lagged variables raised to certain powers, indicates that nonlinearity is essential in forecasting. The use of trigonometric terms is compatible with short-term cyclical behavior that is also essential in exposing additional nonlinearity that is capturable through the good approximation properties of (truncated) Fourier expansions. Finally, the use of transition functions with α coefficient close to unity means that nonlinearity exists also in the intermediate period between smooth transition and abrupt breaks. Specifically, coefficients close to 0.90 probably imply that we have structural breaks, but coefficients close to 0.80 probably imply that there is a relatively fast regime change and that, therefore, the market is between transition and structural breaks. This complex behavior cannot be captured using simple time-series models that allow structural breaks (at unknown dates), or exclusively smooth transition models that ignore other forms of nonlinearity, like lagged powers of certain variables and / or Fourier terms. Our approach allows for an examination of such complicated patterns using a simple and systematic procedure that can be adapted easily for use in other contexts.

The nonlinear character of several variables that lead to their exclusion from the final model is apparent in Table 7. Further, the fact that many α_c are close to unity indicates structural changes rather than smooth transitions. The model performs equally well or even better relative to what we have presented so far, while its marginal likelihood is slightly better in comparison to what we presented in Table 6. However, we believe that even simpler models, without accounting for breaks or smooth transitions, can perform equally well in the context of BCR forecasting.

Concerning economic interpretation, it appears from Table 7 that the forecasting of the BDI involves variables like HAMNO and HAMDWT, reflecting the number and capacity respectively of the Handymax fleet. Handymax is the largest fleet in the sector (3,303 vessels) and has the second largest capacity in dry bulk after the Panamax fleet. It usually carries grains and fertilizer. We have also used the BCI that reflects demand and supply for capsize services, which are primarily carrying iron ore with annual trade growth in the area of 5%, while the capsize fleet grows at -0.1%. Similarly, the BPI, which represents the second largest fleet in the dry bulk sector, with trade growth for its cargo about 4.6% and fleet growth about 2%. So, supply-and-demand conditions pertaining to the two largest categories of the dry bulk sector, and specifically iron ore trade, seem to play a decisive role in explaining and predicting the BDI. The year 2018 will most likely prove a turning point for the BDI; this is supported by the fact that demand growth outweighs supply consistently for over a year now.

(Helpful insights for this section, 5.2, were drawn from Bunn et al. (1982) and Nembhard et al., 2001).

6. Concluding Remarks

This paper contributes to the existing forecasting literature of the shipping sector. It explores the time path of BDI by proposing a new model based on a set of lagged predictors, as well as lagged values of the BDI itself. It extends previous research that treats BDI mostly as a forward-looking indicator of economic activity, by providing a reliable time path for it on which future prediction of economic activity reflecting global trade may be based.

We have implemented two novel procedures in our study in the context of the Bayesian Compressed Regression. We included transition functions in our predictive set, in order to capture sudden changes in the BDI; then, instead of estimating the transition functions, we let the coefficients vary and searched for the variables in the predictive set that contributed to the most accurate forecasts. Furthermore, we have advanced the literature aimed at predicting BDI, by demonstrating a more

successful out-of-sample forecasting performance, able of capturing the abrupt fall of 2008 but also the rest of behavior of the BDI until present.

Our results provide justification for the use of Bayesian analysis in forecasting and demonstrate that through its proper implementation, forecasting in shipping can be successful. Thus, the main policy implication of our study for shipping is that future performance can be approximated, and consequently result to optimal choice of strategies by stakeholders in the sector.

Future research on forecasting in the shipping industry should extend the scope of studying the time path of the BDI to include sub-indices such as BCI (Baltic Cape Index), BPI (Baltic Panamax Index), BSI (Baltic Supramax Index) and HSI (Baltic Handysize Index). The latter is known to reflect the cost of shipping and hence activity in the four dry bulk markets, Capes, Panamax, Supramax and Handysize.

References

Adland, R. and Cullinane, K. (2005). A Time-Varying Risk Premium in the Term Structure of Bulk Shipping Freight Rates. *Journal of Transport Economics & Policy,* 39(2), 191-208.

Alizadeh, A., Ådland, R. and Koekebakker, S. (2007). Predictive power and unbiasedness of implied forward charter rates. *Journal of Forecasting*, 26(6), 385-403.

Baltic Exchange. Market Report. (2016).
http://www.balticexchange.com/market-information/reports/general

Batchelor, R., Alizadeh, A. and Visvikis, I. (2007). Forecasting spot and forward prices in the international freight market. *International Journal of Forecasting*, 23(1), 101-114.

Beenstock, M. and Vergottis, A. (1989a). An econometric model of the world shipping market for Dry Cargo, Freight and Shipping. *Applied Economics* 21, 339–356.

Beenstock, M., and Vergottis, A. (1989b). An econometric model of the world tanker market. *Journal of Transport Economics and Policy,* 23, 263–280.

Boonstra, H.J. (2005). *Model-based estimation of a finite population total: a Bayesian approach*. Statistics Netherlands: Heerlen.

Bowden, A., Hurburt, K., Aloyo, E. and Marts, C. (2010). *The Economic Costs of Maritime Piracy*. Working Paper: One Earth Future Foundation.

Bunn, D. W. and Kappos, E. (1982). The Use of Bayesian Forecasting to Make Process Adjustments During Transitions. *European Journal of Operational Research*, 9 (2), 173-180.

Carter, R. A. L. and Zellner, A. (2002). *The ARAR error model for univariate time series and distributed lag models (Tech. Rep.).* Department of Economics, University of Western Ontario.

Clarkson Research Database. (2016). https://www.clarksons.net/portal

Cullinane, K. (1992). A short adaptive forecasting modal for BIFFEX speculation: a Box-Jenkins approach. *Maritime Policy and Management*, 2, 91–114.

Cullinane, K. (1995). A portfolio analysis of market investments in dry bulk shipping. *Transportation Research Part B: Methodological*, 29 (3), 181-200.

Cullinane, K., Mason, K. and Cape, M. (1999). A comparison of models for forecasting the Baltie Freight Index: Box-Jenkins revisited. *International Journal of Maritime Economies*, 1(2), 15–39.

Driehuis, W. (1970). An Econometric Analysis of Liner Freight Rates. *Review of Word Economics,* 104 (1), 96-119.

Hellenic Shipping News Worldwide. Dry Bulk Shipping: Poor freight rates despite strong demand growth from China. (2016). http://www.hellenicshippingnews.com

Geman, H. and Smith W.O. (2012). Shipping Market and Freight Rates: An Analysis of the Baltic Dry Index. *The Journal of Alternative Investments*, 15 (1), 98-108

Geweke, J. and Amisano, G. (2010). Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting*, 26 (2), 216-230.

Guhaniyogi, R., and Dunson, D. (2015). Bayesian compressed regression. *Journal of the American Statistical Association*, 110, 1500-1514.

Harlaftis, G., Tenold,S., and Valdaliso,J. (Eds.). (2012). *The World's Key Industry. History and Economics of International Shipping*. London: Palgrave/MacMillan.

ICS, International Chamber of Shipping. World Seaborne Trade. (2016).
ICS-Shipping.org

Kadiyala, K. R. and Sune, K. 1997). Numerical Methods for Estimation and Inference in Bayesian VAR-Models. *Journal of Applied Econometrics*, 12 (2), 99-132.

Kalouptsidi, M. (2014). Time to Build and Fluctuations in Bulk Shipping. *American Economic Review*, 104 (2), 564-608.

Kavussanos, M., and Nomikos, N. (1999). The forward pricing function of the shipping freight futures market. *Journal of Futures Markets*, 19 (3), 353-376.

Kavussanos, M. G., and Alizadeh, A. H. (2001). Seasonality patterns in dry bulk shipping spot and time charter freight rates. *Transportation Research*: Part E, 37 (6), 443-467.

Koekebakker,S., Adland,R. and Sodal, S. (2006). Are Spot Freight Rates Stationary? *Journal of Transport Economics and Policy*, 40 (3), 449-472.

Koop, G., and Korobilis, D. (2013). Large Time-Varying Parameter VARs. *Journal of Econometrics,* 177 (2), 185-198.

Koop, G., Korobilis, D. and Pettenuzzo, D. (2016). *Bayesian Compressed Vector Autoregressions*. Working Paper, Business School-Economics: University of Glasgow.

Lewis, S.M. and Raftery, A.E. (1997). Estimating Bayes factors via posterior simulation with the Laplace-Metropolis estimator. *Journal of the American Statistical Association*, 92, 648–655.

Lin, Y., and Wang, C. (2014). The Dynamic Analysis of Baltic Exchange Dry Index. *International Mathematical Forum*, 9 (17), 803-822.

Makridakis, S. et al. (1986). The forecasting accuracy of major time series methods. *International Journal of Forecasting*, 2(1), 119 - 121.

Makridakis, S., Wheelwright, S.C. and Hyndman, J.R. (1998). *Forecasting: Methods and Applications*. New York: John Wiley and Sons.

Nembhard, H. and Nembhard, D.A. (2001). Synthesis or Selection of Forecasting Models. *European Journal of Operational Research*, 130 (2), 437-448.

Odom, P. (2010). Shipping Indices Signal Global Economic Trends. Globalization and Monetary Policy Institute Annual Report, *Federal Reserve Bank of Dallas*, 28-35.

OECD FACTBOOK, (2016). Share of International Trade in GDP, *Economic, Environmental and Social Statistics*, 66-67.

Oomen, J. (2012) *The Baltic Dry Index: A Predictor of Stock Market Returns?* Master Thesis, Tilburg University: Tilburg.

Papailias, F., Thomakos, D. and Liu, J. (2016). The Baltic Dry Index: cyclicalities, forecasting and hedging strategies. *Empirical Economics*. DOI: 10.1007/s00181-016-1081-9

Parzen, E. (1982). Models for Time Series Analysis and Forecasting. *Journal of Forecasting,* 1(1), 67-82.

Randers, J., and Goluke, U. (2007). Forecasting Turning Points in Shipping Freight Rates: Lessons From 30 Years of Practical Effort. *System Dynamics Review*, 23 (2-3), 253-   284.

Sengupta, R. and Tam, Y. (2009). Recent Movements in the Dry Bulk Index. *Economic SYNOPSES,* Federal reserve banks of St. Louis, 12, 1-2.

Silverberg, A. and Nichita, M. (2009). The Lighthouse in a Sea of Indices: The Baltic Dry Index and its Significance. *Ivey Business Review, 1,* 24-25.

Stopford, M. (2009). *Maritime economics*. London: Routledge.

Terasvirta, T. (1998). *Modelling economic relationships with smooth transition regression*. Handbook of Applied Economic Statistics. New York: Marcel Dekker.

Tsay, R.S. (1998). Testing and Modeling Multivariate Threshold Models. *Journal of the American Statistical Association*, 93, 1188-1202.

World Bank Open Data. Trade's Share in Global GDP. (2016).
https://data.worldbank.org/indicator/NE.TRD.GNFS.ZS

Zellner, A. (1971). *An introduction to Bayesian inference in econometrics*. New York: Wiley.

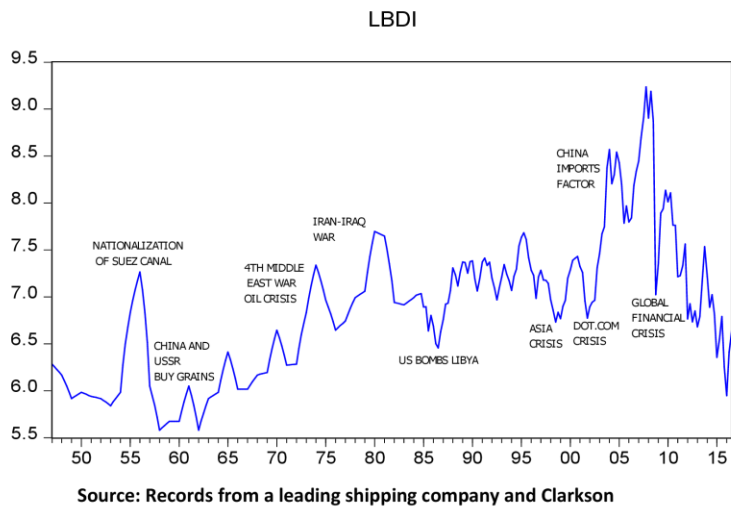**Figure 1: Natural Logarithms of BDI: 1947-2016**

LBDI

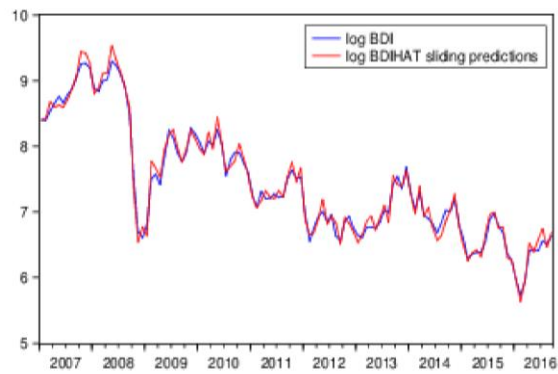Figure 1. Natural Logarithms of BDI: 1947-2016

Figure 2. Forecasting performance of model E for BDI (Sliding Predictions 2007:1 to 2016:6)
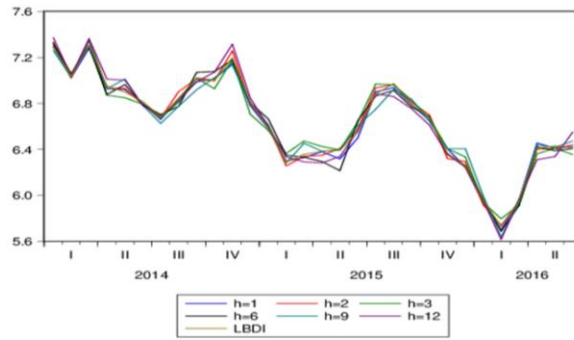
Figure 3. Out of sample Forecasting performance of model E for BDI (2014:1 to 2016:6)
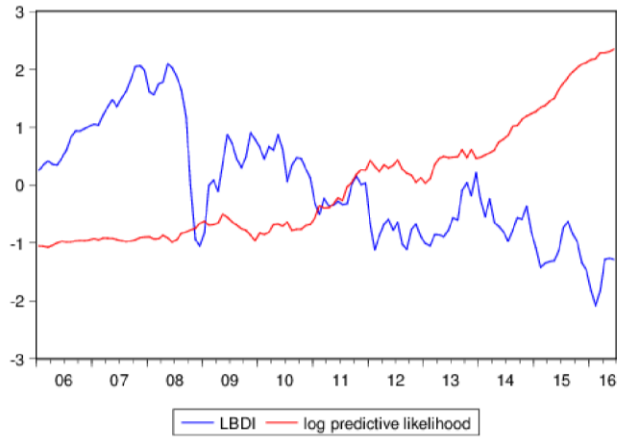
Figure 4. Log predictive likelihood (2006:1 to 2016:6)

Table 1. Data: Variables Used, descriptive statistics and correlations

Panel A: Definition of variables

BDI	BALTIC DRY INDEX
BCI	BALTIC CAPE INDEX
BPI	BALTIC PANAMAX INDEX
HIS	HANDYSIZE INDEX
HAMNO	HANDYMAX FLEET IN NUMBERS
HAMDWT	HANDYMAX IN MILLION DWT
HANDNO	HANDYSIZE FLEET IN NUMBERS
HANDWT	HANDYSIZE FLEET IN MILLION DWT
IRONORE	TOTAL IMPORTS IN THOUSAND TONNES

Panel B: Descriptive Statistics

	BDI	BPI	BCI	HAMDWT	HAMNO	HANDNO	HANDWT	HSI	SHP	IRONORE
Mean	2096.9	2350.5	3832.1	74.6	1509.8	3024.9	80.4	942.2	144.7	49481.5
Median	1449.0	1582.9	2765.9	55.9	1204.0	3069.0	80.8	678.1	120.3	36990.2
Maxim	10843.6	10631.2	16808.1	185.5	3403.0	3350.0	93.2	3221.2	499.6	123173.0
Min	306.9	324.2	868.8	30.4	658.0	2678.0	70.7	197.2	56.9	0.000000
Std.Dev.	1823.6	2078.8	3185.0	46.3	828.8	199.2	6.3	725.9	86.2	29146.6

Panel C: Correlation Matrix

	BDI	BPI	BCI	HAMDWT	HAMNO	HANDNO	HANDWT	HSI	SHP
BDI									
BPI	0.99								
BCI	0.99	0.97							
HAMDWT	-0.65	-0.65	-0.61						
HAMNO	-0.66	-0.65	-0.61	0.99					
HANDNO	-0.68	-0.66	-0.64	0.93	0.94				
HANDWT	-0.67	-0.66	-0.62	0.97	0.97	0.98			
HSI	0.98	0.98	0.95	-0.63	-0.64	-0.65	-0.64		
SHP	0.95	0.94	0.93	-0.64	-0.64	-0.63	-0.63	0.96	
IRONORE	-0.55	0.54	-0.51	0.90	0.90	0.85	0.87	-0.53	-0.55

Table 2: Out-of-Sample Forecasting Performance, MSFEs Relative to Benchmark AR(1) [Using equation (14)]

Predictive Set	Forecast Horizon (months)					
	h=1	h=2	h=3	h=6	h=9	h=12
Only lags of basic set	0.083	1.011	1.133	1.342	1.515	1.617
B. Lags of basic sets, squares and interactions	0.081	0.089	1.103	1.213	1.366	1.891
C. As in B plus trigonometric terms. (K=5)	0.082	1.277	1.381	1.417	1.617	1.781
D. As in B plus transition functions	0.040	0.042	0.051	0.059	0.072	1.025
E. As in C plus transition functions	0.011	0.014	0.019	0.022	0.029	0.035

Table 3: Log Predictive Likelihood Differentials Relative to Benchmark AR(1) (Positive values that the model produces better forecasts than the benchmark).

Predictive Set	Forecast Horizon (months)					
	h=1	h=2	h=3	h=6	h=9	h=12
A. Only lags of basic set	0.0032	0.0015	-0.032	-0.065	-0.082	-0.093
B. Lags of basic sets, squares and interactions	0.0044	0.0023	-0.043	-0.071	-0.091	-0.095
C. As in B plus trigonometric terms. (K=1,…,4)	-0.015	-0.022	-0.045	-0.063	-0.070	-0.085
D. As in B plus transition functions	0.0032	0.015	0.022	0.033	0.055	0.072
E. As in C plus transition functions	0.252	0.257	0.260	0.262	0.268	0.277

Table 4: Out-of-Sample Sensitivity analysis to hyperparameters, maximum absolute percentage change relative to results in Table 2

Predictive Set	Forecast Horizon (months)					
	h=1	h=2	h=3	h=6	h=9	h=12
Only lags of basic set	0.014%	0.017	0.017	0.014	0.014	0.013
B. Lags of basic sets, squares and interactions	0.011%	0.012	0.012	0.013	0.013	0.013
C. As in B plus trigonometric terms. (K=5)	0.014%	0.015	0.011	0.011	0.011	0.011
D. As in B plus transition functions	0.012%	0.012	0.012	0.012	0.011	0.010
E.As in C plus transition functions	0.010%	0.010	0.011	0.010	0.008	0.008

(all numbers in the Table are percentages)

Table 5: Log Predictive Likelihood Differential Relative to Benchmark AR(1), Sensitivity analysis to hyperparameters, maximum absolute percentage change relative to results in Table 3

Predictive Set	Forecast Horizon (months)					
	h=1	h=2	h=3	h=6	h=9	h=12
B. Only lags of basic set	0.003%	0.005	0.005	0.006	0.004	0.004
B. Lags of basic sets, squares and interactions	0.012%	0.007	0.007	0.005	0.003	0.002
C. As in B plus trigonometric terms. (K=1,…,4)	0.004%	0.005	0.003	0.003	0.002	0.004
D. As in B plus transition functions	0.001%	0.003	0.005	0.006	0.004	0.002
E. As in C plus transition functions	0.003%	0.005	0.002	0.003	0.002	0.002

(all numbers in the Table are in percentages)

Table 6: Comparison of MAPE (Mean Absolute Predictive Error) with Alternative Models

Model	MAPE
This Study	1.000
AR(1)	1.414
ARMA(1,1)	1.401
ARIMA(p,d,q)	1.389
ARAR	1.353
ARMAAR	1.355
DARAR	2.441
DARMAAR	2.301
BVAR(2)	1.474
BVAR(4)	1.512
BVAR(6)	1.503
TVP-BVAR(2)	1.381
TVP-BVAR(4)	1.477
TVP-BVAR(6)	1.515

Notes: A number greater than one indicates that the model performs worse relative to the Bayesian Compression model in this study.

Table 7. List of best predictors for BDI

	lags	lags^2	lags^3	Interaction of lags	sin(kπz)	cos(kπz)	transition function
BDI	1-3	1	-	(1,2), (1,3)	z=lag1, k=1 z=lag2, k=1 &2	z=lag1, k=1	lag1, α_c=0.90
BCI	1-4	1	1	-	z=lag1, k=1	z=lag1, k=1	lag1, α_c=0.80, lag2, α_c=0.80
BPI	2-4	1-2	1	-	z=lag1, k=1, z=lag2, k=1 &2	z=lag1, k=1	lag1, α_c=0.90, lag2, α_c=0.90
HSI	-	-	-	-	-	-	-
HAMNO	-	-	-	(2,3)	-	-	-
HAMDWT	-	-	1	(1,2)	-	-	
HANDNO	-	1-2	-	-	-	-	-
HANDWT	-	1	1	-	-	-	-
IRONORE	1-2	1	1	(1,2), (1,3)	-	-	lag1, α_c=0.90, lag2, α_c=0.90