
Article:
Molaris, V.A., Triantafyllopoulos, K., Papadakis, G. et al. (2021) The effect of COVID-19 on minor dry bulk shipping : a Bayesian time series and a neural networks approach. Communications in Statistics: Case Studies, Data Analysis and Applications, 7 (4). pp. 624-638.

https://doi.org/10.1080/23737484.2021.1979434

The Effect of COVID-19 on Minor Dry Bulk Shipping: A Bayesian Time Series and a Neural Networks Approach

V.A. Molaris*[1], K. Triantafyllopoulos†[2], G. Papadakis‡[3], P. Economou§[4], and S. Bersimis¶[5]

[1]Empire Group of Companies
[2]School of Mathematics and Statistics, University of Sheffield, UK
[3]Department of Informatics and Telecommunications, National and Kapodistrian University of Athens
[4]Department of Civil Engineering, University of Patras, Greece
[5]Department of Statistics and Insurance Science, University of Piraeus, Greece

Abstract

Supply chain is a crucial part of the world economy and everyday life. During the outbreak of COVID-19 in 2020, governments across the globe have tried to maintain the supply chain operations in most of the essential goods and services. Dry bulk shipping, a sector characterised by trading versatility and a multitude of cargoes carried, has an important role to play towards that aim. Port restrictions, placed to limit the spread of COVID-19, have created the impression that the market of dry bulk shipping has, to some degree, been affected and become less predictable, as many other transporting services. In this paper, this belief is investigated using both state space modeling and neural networks. In particular, using these methods, one month ahead predictions (during 2020) of the overall freight cost are obtained for three types of vessels (representing the majority of the minor bulks industry). A good forecast performance is observed and this validates the models and the estimation methods proposed. Our results suggests that there is no significant effect of COVID-19 in these type of vessels and their operations. This can provide useful information to shipping managers and policy makers.

keywords: time series analysis, neural networks, dry bulk shipping

*vmolaris@group-empire.com
†k.triantafyllopoulos@sheffield.ac.uk
‡gpapadis@di.uoa.gr
§peconom@uptras.gr
¶sbersim@unipi.gr

1 Introduction

The economic impact of COVID-19 has affected many industries and services worldwide. The severity of this pandemic can be observed by the Global GDP slowdown, the decline of oil prices, supply chain disruptions (Pak et al., 2020; Jefferson, 2020; Sharif et al., 2020; Ibn-Mohammed et al., 2021). The impact of COVID-19 on world economics is also demonstrated by looking at the trading activity of the top 15 countries (Vidya and Prabheesh, 2020). The study shows that the trading density during the first quarter of 2020, when compared to 2018, has declined from 0.833 to 0.429. The contraction caused in trading activity due to COVID-19 is evident in every aspect of world economy.

On March 2020, port restrictions were placed to limit the spread of COVID-19. Each country was asked to provide and implement its own actions and policies in an effort to limit the widespread of the pandemic (Sabat et al., 2020; Han et al., 2020). The types of actions ranged from border closures for crew changes to more than 14 days in transit restrictions. As a result, a study conducted by Millefiori et al. (2020) demonstrated that during COVID-19, dry bulk maritime mobility has fluctuated between -3.32% and 2.28%. Since shipping accounts for more than 70% of global trade, even the slightest mishap can lead to negative economic outcomes and considerable losses (Brussels: European Community Shipowners' Association, 2017). China imports account for more than 30% of commodities transferred via dry bulk ships (Paris, 2020), placing China at the center of shipping's economic well-being. In November 2019, the Organization of Economic Co-operation and Development (OECD) forecasted a 0.7% decline on Chinese GDP for 2020. During COVID-19, the GDP of Chinas fell by 2.16%, which is three times larger than the initial forecast (OECD, 2020). Consequently, the dry bulk markets took a turn for the worse.

The Baltic Dry Index (BDI) measures trading activity by monitoring the freight cost of moving commodities by ships. Similar to the ship market, BDI is divided into 4 categories, with each one corresponding to a particular ship size. On the 10th of February 2020, BDI dropped to 411 points, marking the day as a 4 year low. The Baltic Capesize Index fell to -254, after turning negative for the first time ever on 31st of January 2020, the Baltic Panamax Index to 546, the Baltic Supramax Index to 484 and the Baltic Handysize Index to 302. Michail and Melas (2020) suggested that Dry bulk shipping was directly affected by the outbreak of COVID-19 - not only by the restrictions imposed globally, but also from the increase in reported COVID-19 cases. This article reports, that an increase of 1% in reported COVID-19 has an adverse effect of 0.03% on BDI. Considering all the above, it is expected that all sizes were affected equally by the pandemic. This paper aims to investigate the behavior of the minor bulks markets, namely Handysizes, Handymaxes and Supramaxes in relation to the COVID-19 pandemic.

Considering the above observations, it might be expected that all sizes were affected equally by the pandemic. However, this is not necessarily the case as e.g. the Handysizes are known to be the most resilient in terms of volatility, earnings and performance when compared to other segments of BDI, seconded

by the Supramaxes. The aim of this paper is to study whether trading activity in the above mentioned three categories of minor bulk industry is affected by COVID-19. In particular, in this study, we analyze appropriate data in order to show that the above mentioned market has not been affected significantly by the global market disruption due to COVID-19 pandemic. Furthermore, in an appropriate unified data set coming from various sources, we apply both Bayesian time series modeling and neural networks in order to explore the appearance or not of a change point on the multivariate time series.

The remainder of the paper is organized as follows. Section 2 gives a description of the data and offers a motivation for the study in Sections 3 and 4. In Section 3, matrix-variate dynamic linear modelling and neural networks are discussed. For the former, we propose a flexible linear model, whose inference follows the Bayesian paradigm, while for the latter we discuss a standard neural network architecture. For both approaches we provide some implementation details. Section 4 gives illustration and commentary on the application of these two methods to the data set. Finally, we give some concluding remarks in Section 5.

2 Motivation and data description

In this section, we describe the motivation for the analysis that follows and we briefly describe the data considered in this paper.

As mentioned earlier, even if the shipping market is in turmoil, worldwide data provide some evidence suggesting that minor bulks trading has not been affected by the COVID-19 pandemic. Motivated by this observation, we propose two statistical models, which can show that trading activity during 2020 does not exhibit systematic changes in comparison to earlier years. This is an important point as inelastic demand can play a central role in the future of shipping strategies worldwide. Dynamic linear models and neural networks for time series are used in order to compare and contrast forecast performance and to demonstrate the effect of COVID-19.

To examine the effect of the COVID-19 pandemic on dry shipping, we merged and unified data regarding contracts for dry cargos from two sources:

1. AXSDry[1], a service offered by ASX Marine under the subscription of Empire Navigation. We gathered all relevant data from January 2001 to November 2020.
2. Clarksons Drycargo Chartering[2], a service offered by Clarksons SIN terminal under the subscription of Empire Navigation. We gathered all relevant data from January 2019 to December 2019.

The collected records are distinguished into two categories, depending on the type of contract, namely:

[1]Under the permission of AXSMarine (It includes data and material from AXSMarine, all rights reserved).

[2]Under the permission of Clarksons Research.

1. *Time charters*. A vessel is hired for a specific period of time, during which the charterer commercially manages the vessel and the owner is responsible for the technical management. More specifically, the charterer undertakes fuel and port costs, whereas the owner undertakes operational expenses, such as overheads, insurance and spares. The charterer pays the owner a daily hire for providing this service.
2. *Voyage charters*. A vessel is hired for a trip from a load port to a discharge port. The charterer pays the owner on a per ton basis, and all expenses are undertaken by the owner.

In the sequel, we focus exclusively on voyage charters, as they provide detailed information about each contract. In contrast, time charters specify only the duration of the hire (usually in months), without any details about the size, the value or the type of products that will be transported. Nevertheless, the voyage charters category still provides a clear insight on the effect of COVID-19 to the dry shipping market.

Before processing the collected voyage charters, noisy or incorrect information was excluded. Firstly, duplicate data, identified as records from different data sources that pertain to the same contract, were removed. Secondly, all records without a specified *laycan* were discarded, where *Laycan* is defined to be the time window that the ship should be ready to load a given cargo. A contract with no laycan record is cancelled. Finally, all records that had a zero or missing hire rate as well as the outliers (defined as contracts whose rates were larger than 3 standard deviations above the mean) were also cleared away.

Among the remaining data, we considered only the contracts that pertain to three types of vessels, depending on their deadweight (in tons):

1. Handysize, with a deadweight in [25,000, 40,000),
2. Handymax, with a deadweight in [40,000, 50,000),
3. Supramax, with a deadweight in [50,000, 60,000).

The resulting data consist of 1,058 records for handysize vessels, 342 records for handymax vessels and 740 records for supramax vessels.

Every record is described by the vessel's deadweight, the laycan period, the hire rate per ton (*rate*) and the size of the cargo in tons (*quantity*). Among these variables, we consider only the vessel type, the rate, the quantity, the value as well as the month and year, as determined by the last date of the laycan. We also define a variable, called *freight cost*, to express the monetary value of a contract, i.e., $freight\ cost = rate \times quantity$. The reason is that each vessel caries different quantity of different goods under different rates. The freight cost aggregates this information and therefore reflects the total monetary value across the vessels.

Due to the temporal sparsity of the collected records, we aggregate all variables according to their vessel type and month id. The values of the latter are $t = 1, 2, \ldots, 240$, as there are 240 months between January, 2001 (id=1) and November, 2020 (id=240). Hence, there are 240 records per vessel type, each

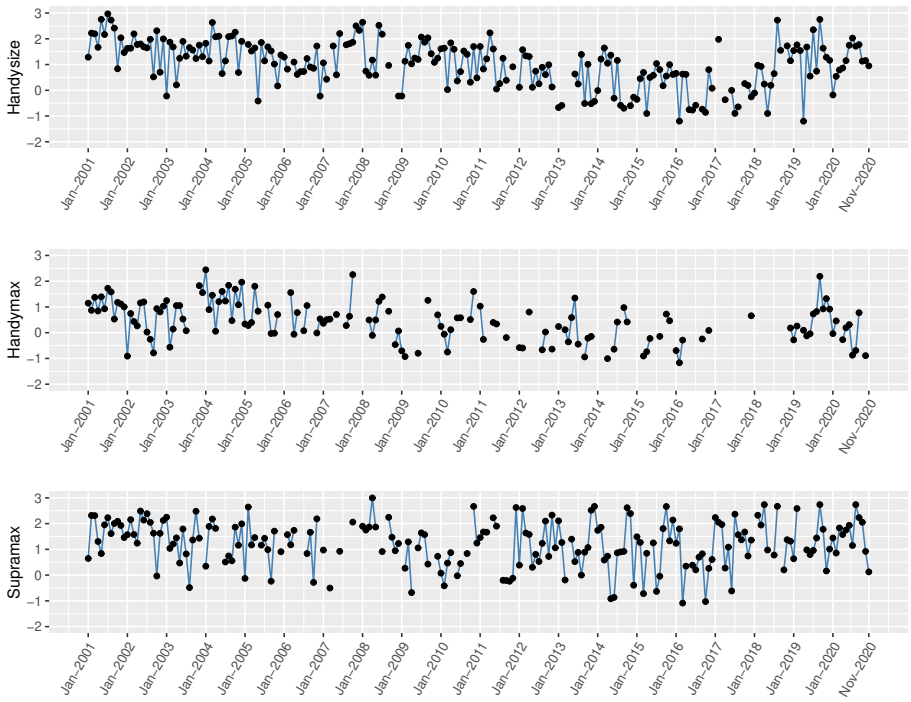

Figure 1: Time series plots for the overall freight cost divided by 10^6 in log scale for the three types of vessels from January, 2001 to November, 2020. Consecutive, non missing observations are jointed together with solid lines.

described by the overall rate, overall quantity and overall freight cost of dry cargo contracts for the corresponding month. Figure 1 presents the time series plots for the overall freight cost, divided by 10^6, in log scale. For each month, we also measure its *frequency*, i.e., the total number of contracts for a particular vessel type. It is noted that minor dry bulks shipping is quite fragmented resulting to scarce data, as indicated in the plots in Figure 1. For this reason, in case of no contracts for a particular month and vessel type, all variables were set to 0 in log scale.

We also consider an external variable, namely the cost of fuel, as determined by their monthly average price in \$/Ton at the Singapore Stock Exchange. From January, 2001 to December, 2019, we consider the High Sulphur Fuel Oil (3.5% Sulphur) as fuel. From January 2020 on, we consider the Very Low Sulphur Fuel Oil (0.5% Sulphur) as fuel. The reason for this transition is the International Maritime Organization (IMO) 2020 regulation[3], where the maximum sulphur content is capped at 0.5%.

[3] `https://www.imo.org/en/MediaCentre/HotTopics/Pages/Sulphur-2020.aspx`

3 Methodology

In this section, we describe the models and methods deployed for the analysis of Section 4. In particular, Sections 3.1 and 3.2 propose time series and neural networks models, respectively, discussing the relevant estimation and forecasting procedures adopted in this paper.

In the following, the overall freight cost $y_t = (y_{1t}, y_{2t}, y_{3t})'$ divided by 10^6 in log scale for each of the three vessel categories is denoted by $z_t = (z_{1t}, z_{2t}, z_{3t})'$, where $z_{it} = \log(y_{it}/10^6)$ at times $t = 1, 2, \ldots, 240$. The response variable corresponds to z_t and y_t for the dynamic linear model and the non-linear model, respectively. At the same time, 10 time-varying variables $x_{it}, i = 1, \ldots, 10$, namely the rate, the quantity and the frequency of handysize, handymax and supramax vessels (3×3) along with the value of oil are included as explanatory covariates.

We use the first 229 months (until December 2019) of each vessel type as a training set for building our prediction models and the remaining 11 months (January-November 2020) to evaluate the forecasts. Although, historical data may not be so reliable as more recent data, it is well known that in time series estimation prior information is deflated over time. As a result, forecasting is not likely to be affected by early observations reliability issues.

3.1 Dynamic linear modelling

3.1.1 Exploratory data analysis

A first observation is that y_{it} exhibit autocorrelation ($y_{it}, y_{i,t-s}$ are correlated, for some s) as well as cross-correlation (y_{it} and y_{jt} are correlated, for $i = 1, 2, 3$). Hence, a suitable multivariate time series model is required, which will enable estimation of these correlations as well as forecasting of y_{it}. A second observation is that the histograms of y_{it} suggest a gamma or a log-normal distribution. Figure 2 shows histograms of the time series y_{it} ($i = 1, 2, 3$) together with their respective autocorrelation (ACF) plots. The histograms, shown in the left panels, indicate a strong skewness as mentioned above. The ACF plots, shown in the right panels, indicate that handysize and handymax (y_{1t}, y_{2t}) are marginally non-stationarity while supramax (y_{3t}) appears to be stationary. The three time series exhibit some cross-correlation. Considering the vector time series y_t, the sample correlation matrix is

$$\rho(\{y_t\}) = \begin{bmatrix} 1 & 0.408 & 0.145 \\ 0.408 & 1 & 0.159 \\ 0.145 & 0.159 & 1 \end{bmatrix}.$$

Clearly handysize and handymax are more corrrelated ($\rho = 0.408$), than handysize with supramax ($\rho = 0.145$) and handysize with supramax ($\rho = 0.159$). We remark, though, that the Pearson correlation coefficient is not a reliable measure of correlation. A more reliable estimation of the correlation matrix can

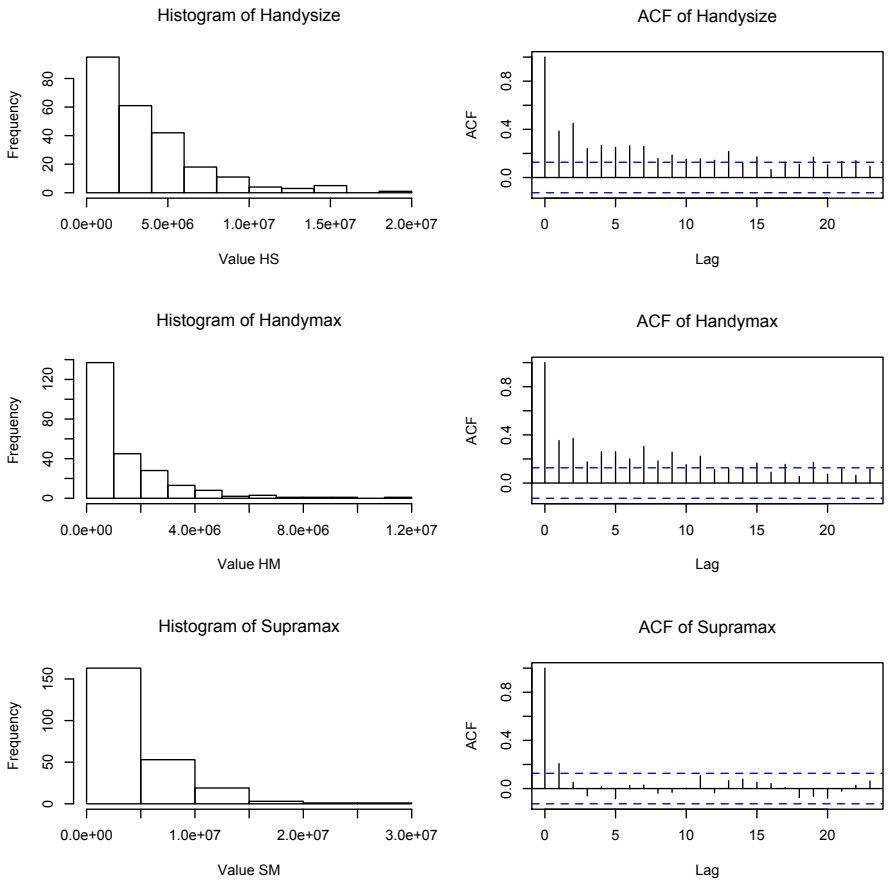

Figure 2: Histograms (left panel) and ACF (right panel) of the time series y_{1t} (handysize), y_{2t} (handymax) and y_{3t} (supramax).

be obtained by the dynamic linear model discussed in the sequel. More specifically, dynamic linear modelling proposes a Bayesian estimation of the covariance matrix of the observations, which is a more accurate method to estimate the correlation among the three time series. For a general description of exploratory data analysis in time series, the reader is referred to Shumway and Stoffer (2017).

3.1.2 Dynamic linear model

In this paper we adopt a Gaussian dynamic linear model for the vector time series $z_t = (z_{1t}, z_{2t}, z_{3t})'$, where $z_{it} = \log(y_{it}/10^6)$. State-space or dynamic linear models, discussed in detail in West and Harrison (1997) or Shumway and Stoffer (2017), are regression-type models, whose parameters follow a Markov

process to describe local evolution. They encompass the popular ARIMA-type or Box-Jenkins models, reported in the classic textbook Box et al. (2008). Here we briefly describe the dynamic linear model of Triantafyllopoulos (2008), which is suitable to enable the estimation of the cross-correlation covariance matrix of z_t. This model consists of of the following observation equation

$$z_t' = x_t'\Theta_t + \epsilon_t', \quad \epsilon_t \sim N(0, \Sigma), \tag{1}$$

where x_t is a time-varying $d \times 1$ design vector, Θ_t is a $d \times 3$ parameter matrix, ϵ_t is a 3×1 white noise vector, which is assumed to follow the trivariate Gaussian distribution with zero mean vector and some covariance matrix Σ, which is subject to estimation. The design vector x_t includes intercept and time-varying covariates. Together with equation (1) this model considers the evolution of the parameter matrix Θ_t according to the following Markov-process

$$\Theta_t = F\Theta_{t-1} + \Omega_t, \tag{2}$$

where F is a $d \times d$ transition matrix and Ω_t is a $d \times d$ innovations sequence, which is assumed to be internally independent and independent of ϵ_t. It is also assumed that Ω_t follows a matrix-variate Gaussian distribution, $\Omega_t \sim N(0, W_t, \Sigma)$, where W_t is a $d \times d$ covariance matrix and Σ is described above. This means that $\text{vec}(\Omega_t)$ follows a $3d$-multivariate Gaussian distribution, i.e. $\text{vec}(\Omega_t) \sim N(0, \Sigma \otimes W_t)$, where $\text{vec}(\cdot)$ is the column stacking operator of a matrix and $\otimes$ denotes the Kronecker product of matrices; matrix-variate distributions are discussed in detail in Gupta and Nagar (2000).

The model is completed by specifying a prior distribution for Θ_0 as $\Theta_0 \sim N(0, C_0, \Sigma)$, for some known covariance matrix C_0. In the following, we assume that the model components F, W_t and the prior moments m_0 and C_0 are known; their specification or estimation is discussed in detail in West and Harrison (1997) and discussed in this paper in Section 3.1.3.

Conditional on Σ, the Kalman filter applies, so that with $z^n = (z_1, \ldots, z_n)$ a collection of observations, the posterior distribution of Θ_t is a matrix-variate Gaussian distribution

$$\Theta_t \mid z^t \sim N(m_t, C_t, \Sigma), \tag{3}$$

where m_t and C_t are provided by the Kalman filter recursions (West and Harrison, 1997), for $t = 1, \ldots, n$. In order to estimate Σ, the covariance matrix of ϵ_t, we adopt a Bayesian framework. We set an inverse Wishart prior for Σ, so that the inverse of Σ follows the Wishart distribution $\Sigma^{-1} \sim W(\nu_0, \nu_0 S_0)$, for some parameters ν_0 and S_0. It follows that the posterior distribution of Σ is inverse Wishart, or $\Sigma \mid z^t \sim W(\nu_t, \nu_t S_t)$, with recursions of nu_t and S_t depending on m_t and C_t from equation (3). Finally, (3) combined with the Wishart posterior of Σ^{-1}, results in the posterior distribution of Θ_t, unconditionally of Σ as a matrix-variate Student t distribution, i.e. $\Theta_t \mid z^t \sim t(nu_t, C_t, S_t)$, for $t = 1, 2, \ldots, n$. Based on this information, the h-step ahead forecast distribution is a trivariate Student t distribution, given by

$$z_{n+h} \mid z^n \sim t(\nu_n - 2, f_{n+h}, Q_{n+h} S_n),$$

where $f'_{n+h} = x'_{n+h} m_n(h)$, $Q_{n+h} = x'_{n+h} C_n(h) x_{n+h} + 1$ and $m_n(h)$, $C_n(h)$ are provided by a repeated application of the transition equation (2) together with the posterior distribution (3). For more details and derivations, the reader is referred to West and Harrison (1997) and Triantafyllopoulos (2008).

3.1.3 Model specification

In this section, we discuss the specification of the model components used in the data analysis, which follow in Section 4. We adopt the dynamic linear model (1)-(2), together with the priors of Θ_0 and Σ discussed above.

The 11×1 design vector x_t includes a unit (corresponding to the time-varying intercept) and 10 time-varying covariates x_{it}, which are x_{1t}, x_{2t}, x_{3t}: the rate, the quantity and the frequency of handymax at time t, x_{4t}, x_{5t}, x_{6t}: the rate, the quantity and the frequency of handysize, x_{7t}, x_{8t}, x_{9t}: the rate, the quantity and the frequency of supramax and x_{10t} the value of oil. The variables x_{2t}, x_{5t}, x_{8t} (quantity) are divided by 10^4. The transition matrix F of (2) is set to be the 11×11 identity matrix, so that the state-matrices Θ_t evolve according to the random walk $\Theta_t = \Theta_{t-1} + \Omega_t$. The data structure of z_t supports a slow evolution of the time-varying coefficient matrix Θ_t, so that $E(\Theta_t \mid z^{t-1}) = E(\Theta_{t-1} \mid z^{t-1})$, but with increased covariance matrix or $\text{Var}[\text{vec}(\Theta_t) \mid z^{t-1}] \geq \text{Var}[\text{vec}(\Theta_{t-1}) \mid z^{t-1}]$, where $\text{vec}(\cdot)$ is the column stacking operator of a matrix as in Section 3.1.2.

An efficient way to model this variance is by using a discount factor $0 < \delta \leq 1$ so that $\text{Var}[vec(\Theta_t) \mid z^{t-1}] = \delta \text{Var}[vec(\Theta_{t-1}) \mid z^{t-1}]$. This implies the specification $W_t = (1-\delta)\delta^{-1} C_{t-1}$, where C_{t-1} is the left covariance matrix of Θ_{t-1} at time $t-1$, see also equation (3). If $\delta = 1$ we obtain a static evolution for $\Theta_t = \Theta$ as $W_t = 0$ in this case. If δ is lower, this setting results in erratic forecasts, exhibiting high variability. After consulting performance measures as well as likelihood optimisation, we have chosen $\delta = 0.98$. This value of δ agrees with suggestions of West and Harrison (1997) and offers stable forecasts, which are able to adapt well to information coming at each point of time.

Finally, following a weakly informative prior approach, we have set the priors as

$$m_0 = 0, \quad C_0 = 1000 I_{11}, \quad \nu_0 = 3 \quad \text{and} \quad S_0 = 3^{-1} I_3,$$

where I_k denotes the $k \times k$ identity matrix ($k = 3, 11$). The prior $m_0 = 0$ essentially sets the forecast at time $t = 1$ equal to the zero vector; however, the model learns from the observations very quickly and this setting does not have an effect in the forecasts considered after the training time. The left covariance matrix C_0 of Θ_0 is set to have a large value (precision matrix $C_0^{-1} \approx 0$) reflecting large uncertainty (or low precision) around our choice for the prior mean $E(\Theta_0) = m_0$. For the covariance matrix Σ, the prior degrees of freedom $\nu_0 = 3$ are chosen so that the Wishart prior distribution of Σ^{-1} is non-singular. Experimentation reveals that $\nu_0 S_0 = I_3$ works best, hence we use $S_0 = 3^{-1} I_3$.

We remark that alternative specification or estimation of the model components may be used. For example we may use the EM algorithm for stat-space

models (Shumway and Stoffer, 2017) or Markov chain Monte Carlo methods (West and Harrison, 1997) in order to estimate the matrices F and $W_t = W$, if W_t is assumed to be time-invariant. In this paper we have decided not to pursue such approaches, because they slow considerably the computational speed of the algorithm, without offering significant gains.

3.2 Non-linear modeling using neural networks

The neural networks, in all their variations, are crafted for classification and regression tasks (Goodfellow et al., 2016). In both cases, they receive as input a vector $x = (x_1, \ldots, x_n)'$, where x_i is a numerical feature that corresponds to the i^{th} independent variable, and produce as output a vector $y = (y_1, y_2, ..., y_m)'$, where y_j denotes the prediction for the j^{th} dependent variable.

In our case, the output vector encompasses the overall freight cost in a particular month t, i.e., $y_t = (y_{1t}, y_{2t}, y_{3t})'$, where y_{1t}, y_{2t} and y_{3t} represent the overall freight cost of Handysize, Handymax and Supramax, respectively, at month $t = 1, 2, \ldots, 240$. The input vector comprises of the values $x_{i(t-1)}, i = 1, \ldots, 10$ of the ten explanatory variables at the previous month $t-1$ along with the corresponding overall freight cost of Handysize, Handymax and Supramax vessel category denoted, for simplicity, as $x_{11(t-1)}, x_{12(t-1)}$ and $x_{13(t-1)}$ respectively.

Among the various types of artificial neural networks, we opted for a Multilayer Perceptron (MLP) Goodfellow et al. (2016), due to the limited size of the dataset. Preliminary experiments revealed that more complex models, like the Long Short-Term Memory Networks Goodfellow et al. (2016), convey no benefits in terms of effectiveness, even though they are suitable for time prediction tasks in theory. The reason is that such deep neural networks are crafted for vast volumes of data that typically comprise several million records.

The effectiveness of MLP depends heavily on its topology and configuration. The size of the input and the output layers are task-specific, being equal to the number of independent and dependent variables, respectively. Hence, we have 13 input and 3 output neurons in our case. For the internal structure, we experiment with two different architectures: (i) a single hidden layer with $k \in [1, 20]$ neurons, and (ii) two single hidden layers, the first with $k_1 \in [1, 20]$ neurons and the second with $k_2 \in [\ \lfloor k_1/3 \rfloor, \lceil 2 \cdot k_1/3 \rceil\ ]$. Considering more layers or neurons per layer would result in overfitting, due to the limited dataset size.

The rest of the configuration parameters were set to standard options widely used in the literature Goodfellow et al. (2016): (a) we selected the rectified linear unit function as the activation function; (b) we opted for Adam with its default learning rate (0.001) for the optimization algorithm. The number of epochs was set to 50, which is sufficient for MLP to converge both during training and testing. The batch size was set to 12, so that it includes the records of a whole year. As the objective function to be minimized during training, we selected the mean squared error. For the implementation of MLP, we used the Keras library[4], version 2.4.0.

[4] `https://keras.io`

Table 1: Topology of the MLP ANN used.

Parameter	Value
number of layers	1
number of inputs	13
number of outputs	3
number of neurons in the hidden layer	8

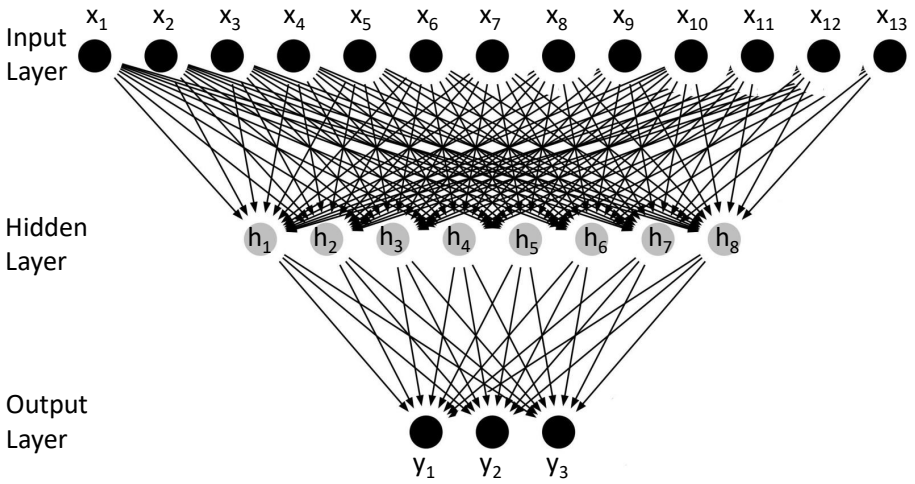

Figure 3: Topology of the MLP ANN used. Every neuron in the input layer (x_1, ..., x_{13}) is connected with every neuron in the hidden layer (h_1,..., h_8) through a weighted, directed edge. Similarly, every neuron in the hidden layer is connected with every neuron in the output layer (y_1, y_2, y_3) through a weighted, directed edge.

Using Grid Search, we assessed the performance of all topologies defined above in order to identify the best one. We observed that MLP is able to exhibit very low error either on Handymax or on Handysize vessels, while the error over Supramax vessels remains consistently high. The topology we selected as optimal minimizes the total root mean squared error, thus achieving a balanced performance across all vessel types. As reported in Table 1 and illustrated in Figure 3, it consists of a single hidden layer with 8 neurons. Due to the stochastic nature of MLP, the predictions we report correspond to the average over 10 runs.

4 One-month ahead forecasts

In Table 2, the predicted values (one-step ahead forecasts) for the period January 2020 to November 2020 are given using both methods presented in the previous

section. Specifically, the first column depicts the overall value of freight cost (in log scale after being divided by 10^6), while the second and the third columns show the predicted values obtained by the dynamic linear model and the neural network, respectively. Furthermore, the fourth and the fifth columns present the lower and upper prediction intervals obtained by the dynamic linear model.

The Mean Absolute Deviation (MAD), the Root Mean Square Error (RMSE), and the Mean Absolute Percentage Error (MAPE) are three of the most frequently used Key Performance Indicators (KPI) to measure, evaluate and compare forecast accuracy. Each one of the three distance measures represents the corresponding distance between the actual and the predicted values and were calculated under the two different methods, i.e. dynamic linear model and the neural network, used to obtain the forecasts over the first 11 months of 2020. The values of these KPIs are also presented in Table 2,.

For Handysize, the neural network exhibits a better forecast performance in comparison to the dynamic linear model. The statistics employed show the cumulative errors are less for the neural network. However, for Handymax and Supramax, the dynamic linear model produces smaller prediction errors. Moreover, it is worth noticing, that the one month ahead prediction of the overall value of freight cost of these three vessel categories, does not seem to be considerably affected by the dynamically changing environment of the COVID-19 outbreak.

Figure 4 shows predicted values of each method together with the actual observations and 95% predictive intervals (produced by the dynamic linear model). It turns out that even for the Handysize category, in which the neural network seems to work better, the dynamic linear model still provides acceptable predictions for all months. On the other hand, the predictions based on the neural network fail to deliver an acceptable prediction for at least one month for the Handymax and the Supramax catregory (January and July respectively). To sum up, it seems that the dynamic linear model has managed to provide better or at least comparative predictions in all cases and for that reason this is the recommended model. Additionally, this model benefits by providing uncertainty analysis via the prediction interval.

Finally, it is worth pointing out the observations of the overall value of freight cost in June, July and October for the Handymax category. For these three months, the actual value was notably smaller, although larger than the lower limit of the prediction interval, than the values predicted by the two methods. The behaviour of Handymaxes in October could be attributed to the second wave of COVID-19 counter measures. However, in June and July a possible explanation could be attributed to the lifting of the restricted measures in relation to the extreme market strategies in a disrupted environment.

5 Conclusions

The analysis conducted in this paper suggests that the Handysize category trading (and economic) activity has not been affected by the COVID-19 pandemic.

Table 2: The actual values (divided by 10^6 in log scale) along with the one month ahead predicted values of the two models for the three vessel categories. The lower and upper predictions intervals based on the dynamic linear model are also included. For each category and method, the three Key Performance Indicators are also reported (smaller values are indicated with **bold** fonts).

		$log(Value/10^6)$	NN Prediction	TS Prediction	TS Lower Forecast	TS Upper forecast
Handysize	January	0.5407	0.1871	0.4397	-0.9115	1.7909
	February	0.7891	0.3602	0.4034	-0.9238	1.7305
	March	0.8642	0.9596	-0.0485	-1.4993	1.4024
	April	1.151	0.8523	0.4559	-0.9577	1.8696
	May	1.7487	1.0577	0.8367	-0.5511	2.2244
	June	2.0282	1.5108	1.7496	0.3573	3.1419
	July	1.7118	1.5003	1.7471	0.2829	3.2112
	August	1.7733	1.8779	1.0543	-0.3859	2.4944
	September	1.1317	1.6309	0.9894	-0.4205	2.3993
	October	1.156	1.2551	0.9093	-0.4202	2.2388
	November	0.9493	0.7885	0.7467	-0.6131	2.1066
		MAD	**0.3146**	0.421		
		RMSE	**0.3684**	0.5244		
		MAPE	**28.1492**	36.1196		
Handymax	January	0.4585	-0.9069	0.245	-0.6326	1.1226
	February	0	-0.2543	-0.1054	-0.9683	0.7574
	March	-0.2718	-0.0549	-0.1264	-1.0684	0.8156
	April	0.1919	-0.2035	-0.0318	-0.9432	0.8796
	May	0.3139	-0.0073	0.0455	-0.8462	0.9371
	June	-0.8732	0.4598	-0.0502	-0.9391	0.8387
	July	-0.6923	0.5225	0.0713	-0.8777	1.0204
	August	0.7759	0.8542	0.4867	-0.4581	1.4314
	September	0	0.5569	-0.2017	-1.1239	0.7204
	October	-0.8906	0.241	-0.0525	-0.9227	0.8177
	November	0	-0.2163	-0.1781	-1.0827	0.7266
		MAD	0.644	**0.3682**		
		RMSE	0.8052	**0.4591**		
		MAPE	104.6592	**58.0058**		
Supramax	January	0.8574	0.8266	0.5	-0.6976	1.6975
	February	1.8309	0.4365	1.2442	0.0661	2.4223
	March	1.5761	1.2337	0.7589	-0.533	2.0509
	April	1.7508	1.0683	1.6609	0.4019	2.9198
	May	1.9334	1.2736	2.014	0.7836	3.2443
	June	1.1499	1.5537	0.7958	-0.4287	2.0203
	July	2.7368	1.2462	2.8962	1.6074	4.1849
	August	2.2346	2.1523	2.0393	0.7713	3.3073
	September	2.0531	1.7877	1.6659	0.4298	2.902
	October	0.9202	1.4807	1.0742	-0.0934	2.2417
	November	0.1248	0.7055	0.5212	-0.6726	1.715
		MAD	0.5903	**0.3253**		
		RMSE	0.7433	**0.3901**		
		MAPE	73.3628	**48.4966**		

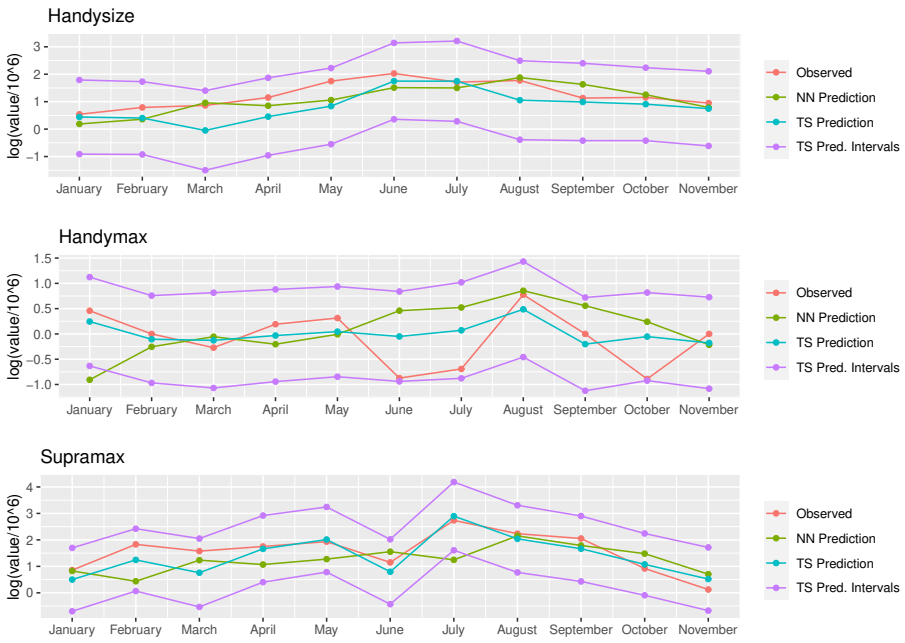

Figure 4: Time series plots of the actual values along with the one month ahead predicted values of the two models. The lower and upper predictions intervals based on the dynamic linear model are also included.

This is also witnessed in case of Supramax category, while for the Handymaxes there seems to be a limited effect for some months. These suggestions may give significant information to shipping strategists for the near future.

A number of factors may justify this performance. The drop in oil prices and subsequently at bunker fuel prices could have a positive impact on these vessels earnings by lowering their voyage costs. The continuous scrapping of old tonnage due to IMO 2022 Water Ballast Management may have limited the supply of available tonnage in the market. In addition, the ordering of new tonnage has decreased 75% as a result of uncertainty caused by the COVID-19 pandemic (Stopford, 2020). The latter statement could create a positive outlook in shipping in terms of supply of ships on an overbuild sector. Finally, due to the world GDP slowdown, big industrial players might have downgraded their parcel size in an effort to limit risk and exposure to long delays at port.

Moreover, the analysis has shown that the performance of the time series based on the dynamic linear models has performed better than the advanced computationally intensive technique of Neural Networks. The dynamic linear model deployed in this paper offers high computational speed and is suitable for on-line estimation and forecasting; its forecast variance and forecast credible bounds also offer a measure of the uncertainty around the point forecasts. On

the other hand, the neural network model offers an advantage in more complex data structures, in particular when the size of the time series is considerably long.

References

G. E. P. Box, G. M. Jenkins, and G. C. Reinsel. *Time Series Analysis: Forecasting and Control*. Wiley, New York, 4th edition, 2008.

Brussels: European Community Shipowners' Association. Shipping and global trade. toward to eu external shipping policy, 2017. URL `https://www.ecsa.eu/images/PositionPapers2017/2017-02-27-ECSA-External-Shipping-Agenda-FINAL.pdf`. Accessed: 2020-12-06.

I. J. Goodfellow, Y. Bengio, and A. C. Courville. *Deep Learning*. Adaptive computation and machine learning. MIT Press, 2016.

A. K. Gupta and D. K. Nagar. *Matrix-Variate Distributions*. Chapman and Hall, New York, 2000.

E. Han, M. M. J. Tan, E. Turk, D. Sridhar, G. Leung, K. Shibuya, N. Asgari, J. Oh, A. Garcia-Basteiro, J. Hanefeld, A. Cook, L. Hsu, Y. Teo, D. Heymann, H. Clark, M. McKee, and H. Legido-Quigley. Lessons learnt from easing covid-19 restrictions: an analysis of countries and regions in asia pacific and europe. *The Lancet*, 396(10261):1525 – 1534, 2020.

T. Ibn-Mohammed, K. B. Mustapha, J. Godsell, Z. Adamu, K. A. Babatunde, D. D. Akintade, A. Acquaye, H. Fujii, M. M. Ndiaye, F. A. Yamoah, and S. C. L. Koh. A critical analysis of the impacts of covid-19 on the global economy and ecosystems and opportunities for circular economy strategies. *Resources, Conservation and Recycling*, 164:105169, 2021.

M. Jefferson. A crude future? covid-19s challenges for oil demand, supply and prices. *Energy Res Soc Sci*, 68:101669, 2020.

N. A. Michail and K. D. Melas. Shipping markets in turmoil: An analysis of the covid-19 outbreak and its implications. *Transportation Research Interdisciplinary Perspectives*, 7:100178, 2020.

L. M. Millefiori, P. Braca, D. Zissis, G. Spiliopoulos, S. Marano, P. K. Willett, and S. Carniel. Covid-19 impact on global maritime mobility, 2020.

OECD. *OECD Economic Outlook, Volume 2020 Issue 2*. 2020. doi: https://doi.org/10.1787/16097408. URL `https://www.oecd-ilibrary.org/economics/oecd-economic-outlook_16097408`.

A. Pak, O. A. Adegboye, A. I. Adekunle, K. M. Rahman, E. S. McBryde, and D. P. Eisen. Economic consequences of the covid-19 outbreak: the need for epidemic preparedness. *Frontiers in Public Health*, 8:241, 2020.

C. Paris. China's import surge drives optimism in dry-bulk shipping demand, 2020. URL https://www.wsj.com/articles/chinas-import-surge-drives-optimism-in-dry-bulk-shipping-demand-11603473327. Accessed: 2020-12-06.

I. Sabat, S. Neuman-Böhme, N. E. Varghese, P. P. Barros, W. Brouwer, J. van Exel, J. Schreyögg, and T. Stargardt. United but divided: Policy responses and people's perceptions in the eu during the covid-19 outbreak. *Health Policy*, 124(9):909 – 918, 2020.

A. Sharif, C. Aloui, and L. Yarovaya. Covid-19 pandemic, oil prices, stock market and policy uncertainty nexus in the us economy: Fresh evidence from the wavelet-based approach. *International Review of Financial Analysis*, 70: 101496, 2020. doi: 10.2139/ssrn.3574699.

R. H. Shumway and D. S. Stoffer. *Time Series Analysis and Its Applications: With R Examples*. Springer, New York, 4th edition, 2017.

M. Stopford. Coronavirus, climate change smart shipping — three maritime scenarios: 2020 – 2050. *Seatrade Maritime News*, pages 1–21, 2020.

K. Triantafyllopoulos. Multivariate stochastic volatility with Bayesian dynamic linear models. *Journal of Statistical Planning and Inference*, 138:1021–1037, 2008.

C. T. Vidya and K. P. Prabheesh. Implications of covid-19 pandemic on the global trade networks. *Emerging Markets Finance and Trade*, 56(10):2408–2421, 2020.

M. West and P. J. Harrison. *Bayesian Forecasting and Dynamic Models*. Springer, New York, 2nd edition, 1997.