MESSAGE IN A BOTTLE: FORECASTING WINE PRICES

Bernardina Algieri, Leonardo Iania, Arturo Leccadito, Giulia Meloni

LIDAM Discussion Paper LFIN
2023 / 04

LFIN

Voie du Roman Pays 34, L1.03.01 B-1348 Louvain-la-Neuve

Tel (32 10) 47 43 04

Email: lidam-library@uclouvain.be

https://uclouvain.be/en/research-institutes/lidam/lfin/publications.html

Message in a Bottle: Forecasting wine prices.*

Bernardina Algieri† Leonardo Iania‡ Arturo Leccadito§

Giulia Meloni¶

June 22, 2023

Abstract

Can we predict fine wine and alcohol prices? Yes, but it depends on the forecasting horizon you choose. We make this point by considering the Liv-ex Fine Wine 100 and 50 Indices, the retail and wholesale alcohol prices in the US for the period going from January 1992 to March 2022. We use rich and multifaceted datasets of economic, survey and financial variables as potential price drivers and adopt several combination/dimension reduction techniques to extract the most relevant determinants. We build a comprehensive set of models and compare forecast performances across different selling levels and alcohol categories. We show that it is possible to predict fine wine prices for horizons of twelve months and retail/wholesale alcohol prices at horizons ranging from one month to two years. Our findings stress the importance of including consumer survey data and macroeconomic factors (international factors and mature market equity risk factors) to sharpen predictions of retail/wholesale (fine wine) prices .

Keywords: *Liv-ex Fine Wine Indices; Alcohol retail and wholesale prices; Reduction data methods; Forecasting Models*

JEL: *C01; C13; C51; C52; E37; G17*

*We thanks Jean-Philippe WEISSKOPF and the participants of the 4th Alliance for Wine & Hospitality Management (ARWHM) Workshop Bolzano 2022 for useful comments. Leonardo Iania acknowledges the support of the following research grant: ARC 18/23-089.

†University of Calabria, Department of Economics, Statistics and Finance, Ponte Bucci, Rende (CS), 87036, Italy & Zentrum für Entwicklungsforschung, ZEF, Universität Bonn, Walter-Flex-Strasse 3, Bonn 53113, Germany. E-mail: b.algieri@unical.it

‡CORE/LFIN, Université catholique de Louvain (UCLouvain), Voie du Roman Pays 34, B-1348 Louvain-la-Neuve, Belgium & University of Leuven (KU Leuven), Department Accounting, Finance and Insurance, Naamsestraat 69, 3001, Leuven, Belgium. E-mail address: iania.leonardo@gmail.com

§University of Calabria, Department of Economics, Statistics and Finance, Ponte Bucci, Rende (CS), 87036, Italy. E-mail: arturo.leccadito@unical.it

¶LICOS Center for Institutions and Economic Performance, Department of Economics, University of Leuven (KU Leuven), Leuven, Belgium & ECARES, Université Libre de Bruxelles, Brussels, Belgium. E-mail address: giulia.meloni@kuleuven.be

1 Introduction

Forecasting wine prices is important for wine producers, sellers, consumers, professionals, and investors, who are interested in knowing future price dynamics to plan production, design buying and selling decisions, outline marketing strategies, form expectations and even invest in alternative classes of assets (Mundi and Kumar, 2023; Bazen and Cardebat, 2022; Dimson et al., 2015). Indeed, pushed by the digitalisation of trade and more sophisticated trading infrastructure, wine has become both a popular consumption good and a newer investment asset that is starting to attract large amounts of funds (Mundi and Kumar, 2023; Ameur et al., 2022; Faye et al., 2015; Fogarty and Sadler, 2014; Masset and Henderson, 2010). It turned out that wine, measured by the Liv-ex Fine Wine 100 index, was a more fruitful investment than several traditional markets in 2021, outperforming the Dow Jones, FTSE 100 and gold (Liv-ex, 2022[1]). Nowadays, wine trading takes place in various venues, including wine exchanges, specialised online stores, traditional or online auctions, trading platforms and over-the-counter (OTC) markets (Oleksy et al., 2021).

The present study aims to forecast fine wine prices - as proxied by the Liv-ex Fine Wine 100 and 50 Indices - plus retail and wholesale alcohol prices in the US and assess the performance of different forecasting models in order to identify the most robust and accurate one. In particular, we focus on horizons ranging from one month to two years and compare a set of forecasting models to a group of benchmarks, including the random walk, and different variants of growth and distributed lag models.

We contribute to the emerging literature on forecasting wine and alcohol prices in several ways. To start with, this is one of the first studies that tries to predict (i) fine wine prices using the Liv-ex Fine Wine 50 and 100 Indexes and both (ii) wholesale and retail alcohol prices. In general, the few forecasting studies on wine prices have mainly analysed the Bordeaux red wine price (Bazen and Cardebat, 2022; Ashenfelter, 2008), while, to the best of our

[1] https://www.liv-ex.com/2022/01/industry-benchmark-closes-2021-23/

knowledge, the literature in wholesale and retail alcohol prices is absent. We explicitly focus on the Liv-ex Fine Wine (50) 100, which tracks the (50) 100 most traded fine wines across the world, including the Bordeaux, the Burgundy, the Champagne, and the Rhône. This index, which gives each wine a weight, is a good gauge of the movement and performance of the wine market. Yeo et al. (2015) have also considered the Liv-ex Fine Wine 100 Index, but differently from their analysis, which is based on machine learning techniques, we use parametric regression models and a comprehensive set of potential predictors. Retail and Wholesale wine and other alcohol beverage prices are for the US and their inclusion in our study allows us to assess the forecasting performances of our explanatory variables across different selling levels and alcohol categories. Second, differently from Bazen and Cardebat (2022, 2021) and Cardebat et al. (2017), we apply large-scale dataset techniques to capture the most relevant group of drivers of wine prices useful to make predictions. In detail, we use four sets of methods, all derived from the big data literature: (i) Lasso, (ii) Elastic Net, (iii) Ridge regression, and (iv) Principal Component Analysis (PCA). To the best of our knowledge, this approach is novel in the wine literature. Third, we include a broad range of explanatory variables and use different data sources, such as: (i) macroeconomic data obtained from McCracken and Ng (2020) who compiled one of the most popular datasets in macroeconomics and finance for the US; (ii) consumer surveys data, obtained from the University of Michigan; (iii) commodity prices indexes, taken from the IMF; (iv) stock market data, retrieved from the Kenneth French website; (v) international indicators of key macroeconomic/trade variables, provided by the Fed; (vi) liquidity variables, and (vii) Equity Market Volatility trackers. Finally, we confront different forecasting models to a set of benchmarks, to assess their accuracy and identify the best-performing ones. No studies have been found comparing performances across several models and specifications.

Our results can be summarized as follows. First, when we compare the different machine learning employed in our analysis, we find that no single approach systematically outperform the other ones. Hence, we reinforce the findings of Elliott and Timmermann (2016), who point out that no single model or forecasting method can be expected to consistently outperform other models in forecasting economic and financial series. Second, fine wine prices are, on average, more difficult to forecast than retail and wholesale prices. This results holds both when benchmark against a model with lagged information about the dependent variable and (to a lesser extend) against a pure random walk setting. Fine wine prices can be forecasted for longer term horizons, two years, and with the help of international and financial variables, while local (US) economic variables and surveys data are more important in predicting retail and wholesale prices both in the very short and medium-long term, i.e. from one month to two years. Finally, combining models (where a model is based on an homogeneous set of predictors) lead to a generalised improvement in the forecasting results, both for fine wines (especially for short-term horizons) and for and for wholesale and retail prices (for all forecasting horizons).

The remainder of the study is organised as follows. Section 2 presents the literature review. Section 3 introduces the adopted methodology. Section 4 presents the considered data. Section 5 displays the empirical results and discusses the findings. Section 6 concludes.

2 Literature Review

There is a vast literature on forecasting commodity prices starting from the seminal work by Fama and French (1987). The analyses on forecasting wine prices are relatively less developed. Differently from other agricultural commodities, wine does not have a futures market[2]. As a result, market participants do not have the information necessary to form expectations from 'futures prices'. Therefore, forecasting future wine prices provides different market operators with relevant information that allows them to form expectations of price movements and increase market efficiency.

[2] Euronext launched a wine futures market, the WineFex in 2001; however, it failed to attract sufficient attention from investors and thus disappeared by the end of 2002.

Generally, the emerging wine economics literature has explored price dynamics and their drivers, but not price forecasting.[3] For instance, Burton and Jacobsen (1999) pointed out that wine prices have a higher variance than equities, and the price dynamics show a typical boom-bust behaviour. This feature is confirmed by the analysis of Fogarty (2006), who documented that returns on wine have a cyclical pattern. In addition, Fogarty (2006) showed that dearer wines realize larger returns and have lower volatility than less expensive wines. Dimson et al. (2015) examined the long-term investment performance of fine wines and found that high-quality vintage wines provide higher returns than bonds, art, and stamps and present a high correlation with equities. Bouri and Roubaud (2016) explored the co-movements between fine wine and stock prices in the UK and found that fine wine is a hedge against movements in UK stocks, but cannot be an effective safe haven during periods of market turmoil. Le Fur et al. (2016) indicated that Bordeaux fine wines were more volatile during the financial crisis and were less volatile in non-crisis periods, and fine French wines had inverse volatility trends compared to fine non-French wines (Australian, Italian, and American). Nahmer (2020) pointed out that investing in fine wine raises risks more than returns, and the relevant investment costs in the fine wine market lessen returns. Cardebat and Jiao (2018) found significant cointegration between emerging markets, especially Asia, and fine wine markets, and also causality from the former to the latter. China, in particular, was one of the main drivers of fine wine prices.

The first forecasting analyses on wine prices go back to the works by Oczkowski (2001) and Ashenfelter (2008). The two authors have adopted hedonic regressions to predict fine wine prices. The idea behind hedonic regression is to divide the price into two components: one representing the value attributed to some intrinsic features of the good (e.g. its quality or rarity), and the other measuring the price appreciation over time. It is thus possible

[3] Over the past decade, a set of empirical analyses tried to answer the question "Does Terroir Matter?" (Gergaud and Ginsburgh, 2008) or "What is the Value of Terroir?" (Cross et al., 2011). These studies and related analyses by Ashenfelter and colleagues find mixed effects: Gergaud and Ginsburgh (2008) and Cross et al. (2011) find no effect of "terroir" (production location) on wine prices, while Ashenfelter and Storchmann (2010), Gergaud et al. (2017) and Haeck et al. (2019) do find a positive effect.

to account clearly for the heterogeneity among the different wines. However, according to Bazen and Cardebat (2022), hedonic regressions are not the most appropriate approach to making predictions for their nature, since they typically involve hefty matrices of regressors, with resulting multicollinearity issues and likely imprecise and erratic index coefficients. Yeo et al. (2015) have forecast the log price changes (i.e. returns) of fine wines using machine-learning methods. Paroissien (2020), drawing on the agricultural prices forecasting literature, developed a forecasting model for bulk Bordeaux wine prices at annual and monthly frequencies. The author included several leading economic indicators of supply and demand as predictors and found that stock levels and quantities produced had the highest predictive power. Masset and Weisskopf (2022) have proposed an empirical model that aims to determine how Bordeaux wine producers should fix their release prices and how these should evolve based on a set of relevant signals (such as the trend on the secondary wine market or the economic and financial environment). Bazen and Cardebat (2022) have used a state-space approach to provide forecasts of the price of generic Bordeaux red wine. The method decomposes the wine series into trend, cycle, seasonal and irregular components, for which some components can be stochastic and others deterministic. The different parameters of the model are estimated using Gaussian state-space methods, involving the Kalman filter and maximum-likelihood estimation. A possible drawback of the study regards the univariate setting used by the authors, which overlooks any potential explicative wine determinants.

In this study, we contribute to the existing literature on wine price forecasting by considering novel methodologies, new data and comparisons across different models.

3 Methodology

In order to predict wine and alcohol prices, we compute the h-period wine returns (r) defined as the difference between the natural logarithm (ln) of the price at time $t+h$ and the one at time t. Formally:

$$r_{t,t+h} = \ln I_{t+h} - \ln I_t, \tag{1}$$

where I_t is the value of the wine or alcohol price index at time t. We use a direct forecasting approach and link $r_{t,t+h}$ to the last p month-on-month growth rates and a set of (lagged) determinants, X:

$$r_{t,t+h} = \beta_0 + \sum_{i=1}^{p} \beta_i r_{t-i,t+1-i} + \sum_{j=1}^{q} \delta_j' X_{t+1-j} + \epsilon_t \tag{2}$$

where ϵ_t is the error term. Note that when in (2) $q = 1$, we use only the contemporaneous value of the vector of explanatory variables.

Assuming that the number of observations is larger than the number of parameters, model (2) can be estimated by Ordinary Least Squares (OLS). However, if the number of observations is not much larger than the number of parameters (say d), the OLS estimates will have large variance resulting in inaccurate predictions of the future observations of the dependent variable. In any case, when the number of explanatory variables used in the model is large, it is likely that some of them are actually not related to the dependent variable. If we include these irrelevant variables, we make the model unnecessarily complex and difficult to interpret. While it is possible to use OLS and perform variable selection (for instance, based on information criteria) with the aim of excluding irrelevant variables and hence using only a subset of the vector of explanatory variables, we adopt alternative approaches based either on (i) shrinkage (or regularization) or (ii) dimension reduction. Contrary to OLS, the first approach uses all the possible explanatory variables, but it constrains the coefficient estimates, effectively shrinking the coefficient estimates towards zero. This has the effect of reducing the estimates' variance. Dimension reduction techniques, instead, are based on computing l different linear combinations of the explanatory variables. These l linear combinations are then employed as predictors in a linear regression model that is estimated by OLS. Typically, l is chosen to be much smaller that d, hence the name dimension reduction. When indeed $l \ll d$, dimension reduction techniques too end up reducing the variance of the fitted coefficients.

In detail, we use four sets of methods, all derived from the big data literature: (i) Lasso, (ii) Elastic Net, (iii) Ridge regression, and (iv) Principal Component Analysis (PCA). The first three are shrinkage methods, the latter is a dimension reduction technique.

The Lasso and Ridge regressions are forms of penalized regressions. They are designed to improve OLS estimates by performing dimension reduction and/or variable selection when dealing with large datasets and possibly correlated regressors. The Elastic Net algorithm combines the penalty elements of both the Lasso and Ridge regression. In a model with y as dependent variable and Z as the k-dimensional vector of regressors, it leads to the following minimization problem:

$$\hat{\gamma} = \arg\min_{\gamma} \left[\sum_{t=1}^{T} (y_t - \beta_0 - Z_t'\beta)^2 + \lambda \sum_{j=1}^{k} \left[(1-\alpha)\beta_j^2 + \alpha|\beta_j| \right] \right]. \tag{3}$$

In (3), $\gamma = (\beta_0, \beta)'$ denotes the vector of parameters (including the intercept) and λ determines the extent of the shrinkage penalty and its optimal value is in practice derived by cross-validation. The parameter $\alpha \in [0, 1]$, instead, negotiates the relative weights of the two penalties: $\alpha = 0$ corresponds to the Ridge regression specification and $\alpha = 1$ to the Lasso specification. In our empirical application, we use $\alpha = 1/2$ in the estimations involving the EN method.

PCA is a tool that aims at reducing the dimensionality of the data while capturing most of the information. In particular, it seeks a low-dimensional representation of a set of variables containing as much as possible of their variation. For instance, the first principal component is the normalized[4] linear combination of the variables for which the variance is largest. For $i \in \{2, \ldots, l\}$, the i-th principal component is obtained as the linear combination with the largest variance from all the linear combinations that are uncorrelated with all the principal

[4] The sum of the squared coefficients of the linear combination, known as loadings, is set equal to one.

components from 1 to $i-1$. In our empirical analysis, the number of principal components is chosen so that they explain at least 70% of the variability of the data.

The accuracy of a given model m is assessed using the Root Mean Square Error (RMSE),

$$\mathrm{RMSE}_m = \sqrt{\frac{1}{W_h}\sum_{t=1}^{W_h}(\hat{r}_{t+h,m} - r_{t+h})^2} \tag{4}$$

where W_h is the number of predictions we make for the horizon h, $\hat{r}_{t+h,m}$ is the predicted return from model m for time $t+h$ conditional on the information up to time t and r_{t+h} is the realized return h months from time t.

We consider also combinations of forecasts based on M different models,

$$\sum_{m=1}^{M} w_m \hat{r}_{t+h,m} \tag{5}$$

where the weights, obtained from a training sample, are calculated as

$$w_m = \frac{\mathrm{RMSE}_m^{-1}}{\sum_{n=1}^{M}\mathrm{RMSE}_n^{-1}}. \tag{6}$$

The training sample is used to derive a number of out-of-sample forecasts, that, in turn, are used in conjunction with the realized values to obtain, for each model m, the loss function RMSE.

For each of the benchmark model we consider, we divide the RMSE of a given model (or combination of forecasts) by the RMSE of the benchmark model. In this way we obtain the Score Ratio (SR)

$$\mathrm{SR}_m = \frac{\mathrm{RMSE}_m}{\mathrm{RMSE}_{\mathrm{BM}}} \tag{7}$$

Models (or combinations of forecasts) for which the measure (7) is less than one have a better forecasting accuracy of the considered benchmark model. In addition to the SR, we report p-values from the Diebold-Mariano test (Diebold and Mariano, 2002) of equal accuracy of the forecast of the model and the benchmark.

4 Data

Our dependent variables are classified in three groups. The first group pertains to the measure of fine wine prices. We use the Liv-ex[5] Fine Wine 50 (LIVX50) and 100 (LIVX100) Index, which are the industry-leading benchmarks for monitoring fine wine prices in the secondary market. The LIVX100 index, as mentioned, tracks the price movements of 100 of the most exclusive fine wines coming from France, Italy, Spain, the USA and Australia that meet the necessary criteria in terms of trading volume and value. The index is calculated using the mid-price – the mid-point between the highest live bid and the lowest live offer on the secondary market. The mid-price given is for 12x75cl trades. LIVX50 tracks the daily price movements of the Bordeaux First Growths. The indices are available from January 1992 to March 2022[6] on a monthly basis, and have been retrieved from Bloomberg. The other two groups of dependent variables include the seasonally adjusted alcohol price indices of retail sales (wine, beer and liquor stores) and merchant wholesalers (except manufacturers' sales branches and offices: wine, beer and distilled alcoholic beverages sales) retrieved from the St. Louis Fed[7]. The dynamics of our four dependent variables plus the Liv-ex Fine Wine Investables Index (LIVXFWIN) are reported in Figure 1, where July 2001 is the base month.

We use seven potential types of explanatory variables: (i) a large-scale set of US-related macroeconomic variables; (ii) a large-scale set of US-related Equity Market Volatility (EMV)

[5]Liv-ex is the acronym for the London International Vintners Exchange, the first global marketplace for wine trading, founded in 2000 by two stockbrokers James Miles and Justin Gibbs.

[6]The Liv-ex 100 (50) index is available on Bloomberg terminal with the ticker LIVX100 (LIVX50) Index. The series dates back to July 2001, before that period LIVXFWIN Index (Liv-ex Fine Wine Investables Index) has been considered given the strong correlation (above 97 per cent) recorded for the two indices between 2001 and 2018. We do not consider LIVXFWIN, since the series is discontinued since 2018.

[7]Alcohol retail and wholesale prices have mnemonic MRTSSM4453USS and S4248SM144SCEN, respectively.

Figure 1: Dependent variables: Wine and Alcohol price indices

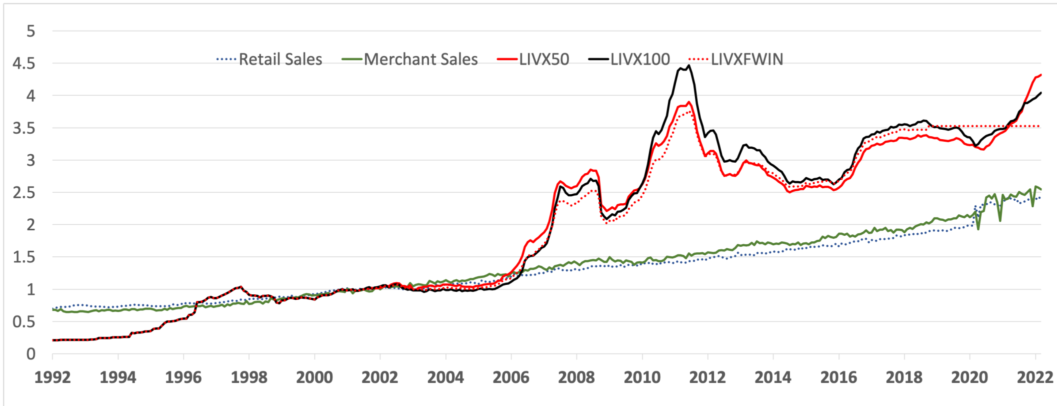

trackers; (iii) surveys of consumer confidence; (iv) international factors related to trade, real activity, exchange rates, and producers/consumers price indexes; (v) stock market risk factors; (vi) liquidity variables; and (vii) commodities prices.

We trace the US-based historical evolution of economic activities via the large-scale dataset of McCracken and Ng (2020), available at the Federal Reserve of St. Louis website. The dataset consists of 126 monthly series classified in 8 groups ranging from financial and nominal variables, such as interest rates and prices, to output data, such as industrial production[8]. The eight groups are listed in Table 1, from row M1 to M8.

Equity Market Volatility (EMV) indexes are 44 newspaper-based series associated with words belonging to three categories, see Baker et al. (2019): (i) economic, economy, and financial conditions (E); (ii) stock market variables (M); and (iii) volatility (V). Masset and Henderson (2010) show that wine returns have low correlations with other assets. Hence, swings in overall equity volatility might be linked to changes in demand (and return) for wine investments, justifying the usage of EMV indices as potential predictor of wine returns. EMV indexes make M9 in Table 1.

Next to historical macroeconomic and uncertainty data, we consider (potential) future

[8] We use the data and data-transformation proposed by McCracken and Ng (2020) at the Federal Reserve of St. Louis website: https://research.stlouisfed.org/econ/mccracken/fred-databases/.

tendencies on economic activity and consumption by using the three major a set of consumers-spending surveys indicators, the Index of Consumer Sentiment, the Index of Consumer Expectations and the Index of Current Conditions, labelled as *ics*, *ice* and *icc* in in Table 1 (see rows M10 to M13). The data is retrieved from the University of Michigan's Survey of Consumers, who conducts this survey on a monthly basis since mid 70's.

The data on international economic indicators have been downloaded from the Database of Global Economic Indicators of the Federal Reserve Bank of Dallas. From this data source we retrieved (i) a measure of inflation, i.e. the headline consumer price index inflation, (ii) a measure of economic activity, namely the output production in the industrial sector, (iii) a trade measure, namely the US dollar value of goods bought from the rest of the world, deflated by the US consumer price index, (iv) an indicator of producers' price, i.e. the changes in the level of prices received by domestic producers for their goods and services, and (v) a nominal exchange rate measure, proxied by the rate at which the US dollar is exchanged for a unit of currency of a bundle of foreign countries. These indexes are obtained as a (US trade) weighted average of a sample of approximately 40 countries, see Martínez-García et al. (2015). In Table 1 these five series are reported from line M14 to M18 as *cpi*, *ip*, *import*, *ppi*, and *ner*, respectively.

As stock market risk factors, we employ the workhorse Fama and French (1993) factors, which are widely known to be most important systematic risk factors in the stock market. The data is freely available at the Kenneth French data library from which we retrieved five factors for the developed markets and five for the emerging economies. Distinguishing between emerging and developing economies might be important in this context as investors seeking for additional returns might turn their attention either to alternative investments such as fine wines or to more riskier but fast markets such as emerging economies. These factors are labeled as *Developed 5 Factors* and *Emerging 5 Factors* in lines M19 and M20 of Table 1.

The sixth set of potential variables, are linked to international (il)liquidity in equity markets. We adopt the Pástor and Stambaugh (2003) measures, divided in aggregated liquidity, liquidity innovation and traded liquidity. Pástor and Stambaugh (2003) find that expected stock returns are related to fluctuations in aggregate liquidity, which account for a sizable proportion of the profits for popular strategy like the momentum strategy. Given the diversification role played by alternative assets (such as wine), changes in stock market liquidity might impact the demand for those assets, and hence their future returns. Liquidity variables make M21 in Table 1. Finally, commodity indexes are obtained from the International Monetary Fund and are listed in appendix A. These indexes are listed as M22 in Table 1.

5 Results

5.1 Setting

We compare the competing models to two baseline settings, based on the random walk and the autoregressive model frameworks:

- BM1: Random walk in growth, i.e. the predicted changes in wine price is equal to its last observed value:

$$\mathbb{E}_t r^w_{t,t+h} = r^w_{t-h,t}. \tag{8}$$

- BM2: A distributed lag model with the last three lagged month on month changes in wine prices as control variables.

$$\mathbb{E}_t r^w_{t,t+h} = \beta_0 + \beta_1 r^w_{t-1,t} + \beta_2 r^w_{t-2,t-1} + \beta_1 r^w_{t-3,t-2}; \tag{9}$$

  which corresponds to (2) with $p = 3$ and $\delta_j = 0 \ \forall j$.

We build-up the alternative forecasting models by augmenting BM2 with explanatory variables obtained via two approaches. The first approach, whose labels begin by 'M', builds-up along two dimensions: the type of data and the dimensionality-reduction method. Table 1 reports

22 models developed along this criteria, whereby (i) we divided the US macro factors along the main classifications (M1 to M8), (ii) we consider a model with the EMV variable (M9), (iii) we group the three Michigan Surveys (M10) and consider them separately (M11 to M13), (iv) we create a model for each international economic factor (M14 to M18), (v) we group developed (M19) and emerging (M20) stock market factors, (vi) we consider a model with the liquidity variable (M21), and (vii) we group the commodities variables (M22). In relation to the dimensionality-reduction criteria, for each model containing more than one variable we apply the three dimensionality-reduction methods outlined in Section 3. It must be noted that when we add only one explanatory variable to BM2, the PCA boils down to standard regression and Lasso/Elastic-net/Ridge are simple penalized regressions. With the second approach, whose labels begin by 'G', we build-up forecasting models by combining the models of Table 1 via the optimal out-of-sample weighting scheme, see eq. (5) and eq. (6) of Section 3. Table 2 reports this set of additional models.

We compare the benchmark and the alternative models via an out-of-sample analysis. We adopt a recursive setting, whereby we produce monthly out-of-sample forecasts for the five benchmark model and all the *Ms* and *Gs* alternative. Each month, the parameters of BM2, *Ms* and *Gs* are obtained using the most recent 15 years of data available at the time of the forecast. The forecasts refer to five forecasting horizons, i.e. $h = 1$, 6, 12, 18 and 24 months, hence covering both short-term (1 month) medium-term (6 months to one year), and long-term (18 and 24 months) holding periods. Using this 15 years rolling window, we make our first prediction for January 2012 ($h = 1$), June 2012 ($h = 6$), December 2012 ($h = 12$), June 2013 ($h = 18$), and December 2013 ($h = 24$). As far as the combined forecasts are concerned, the size of the training set is 24 months. This means that the first prediction we make using this method is for January 2014 ($h = 1$), June 2014 ($h = 6$), December 2014 ($h = 12$), June 2015 ($h = 18$), and December 2015 ($h = 24$). In the end, we obtain two sets of out-of-sample forecasts:

$$\hat{R}_{h,b} = \{\hat{r}_{t+h,b}, \hat{r}_{t+h+1,b}, ...\hat{r}_{T,b}\}, \tag{10a}$$

$$\hat{R}^{rd}_{h,m,p,q} = \{\hat{r}^{rd}_{t+h,m,p,q}, \hat{r}^{rd}_{t+h+1,m,p,q}, ...\hat{r}^{rd}_{T,m,p,q}\}, \tag{10b}$$

where (i) in (10a), which relates to the benchmark forecasts, $b = BM1, BM2$, (ii) in (10b), which refers to the models forecasts, $p = 1, 3$, $q = 1, 3$, $m = M1, ..., M22, G1, ..., G9$, and rd is one of the four dimensionality-reduction method, and (iii) h are the forecasting horizons. The forecasts of (10a) and (10b) are then fed in (4) and (7) to evaluate the (relative) models accuracy.

5.2 Discussion

Overall, for each forecasting horizon and for each dependent variable, we produce a set of 352 alternative predictions[9] using eq. (2), implying that in total we perform 704 models comparisons ($\dim(b) = 2$). A summary of the results for all the methods against BM2, which is the can be found in Table **??**, which compares, for each forecasting horizon, the estimation method/choice of lags in eq. (2), and dependent variable, the percentage of models (or combination of models) that have a RMSE smaller than benchmark model BM2.

Overall, Table 3 points out two main conclusions. First, while PCA seems to marginally be the most effective aggregation method, there is no single model or forecasting method that clearly consistently dominates the other. Except in some specific cases, like the short-term forecast of wholesale, retail and LIVX100 prices, the difference in performance between the two best methods is always below 10%. Hence, our results reinforce the findings of Elliott and Timmermann (2016), who highlight that in forecasting economic and financial series, no single model or forecasting method can be expected to dominate over time. Second, on average, the models are more successful in beating the BM2 when forecasting alcohol retail and wholesale prices rather than fine wine prices. For example, except for forecasting wholesale prices in six

[9] $\dim(p) \times \dim(q) \times \dim(m) \times \dim(rd) = 2 \times 2 \times 22 \times 4$

and twelve months, the best aggregation method outperforms BM2 in at least one case out of two. Potentially, this might be due to the fact that fine wine prices can be considered as an (alternative) investment class, hence as many financial assets, are more difficult to predict, at least in the very short term. Indeed, for the LIVX50 and LIVX100 indexes, the forecasting power of the alternative models is significantly better for two-year horizons.

To have more insights on each model performance, we visualise our large set of results with the help of heatmaps, whereby score ratios smaller (larger) than one are depicted on a red (green) scale which increases in intensity as ratios get smaller (larger). Here we report only the heatmaps obtained using the Lasso method. The heatmaps for the other methods can be found in an online appendix. The heatmaps are depicted in Figure 2 to 17 and refer to the four wine and alcohol indices we consider. In each figure, the five subplots correspond to the forecasting horizons (h), ranging from one month ($h = 1$, top panel) to one year ($h = 24$, bottom panel). In the first panel group (Figures 2-9), the horizontal axes refer to the different explanatory variables (Ms models) we use to characterise the forecasting equation (2) with the Lasso dimensionality reduction method. The twenty-two Ms models for each method are compared to the two benchmark models (BM1 and BM2). For all the panels, the vertical axes refer to the different lags of wine's returns (p) and explanatory variables (q).

Focusing on the performance towards the pure Random Walk model, BM1, a casual inspection of Figures 2-5 confirms the (second) general pattern discovered for BM2 in the summary table: while the alternative models work well in forecasting alcohol retail and wholesale prices, fine wine prices are, on average, more difficult to forecast. The results related to retail and wholesale prices are in line with the general findings of Faust and Wright (2013), which list the Random Walk among the worst model for forecasting CPI inflation. Turning to fine wine, about one model out of four (6 vs 22) does not consistently outperform the BM1 across all horizons. The darkest red cells (the best forecasting performances) are found for longer-term forecasts. This indicates that investors seeking to invest in fine wine products for medium-term horizons might find it valuable to look beyond the past performances of these investment vehicles and consider signals from (i) international macroeconomic developments, trade patterns and commodities prices, (ii) monetary and price data from the largest world economy, and (iii) stock market performances of mature financial markets such as the US and the largest developed economies. Our results on the failure to outperform the Random Walk model for short-term forecasting contrast those of studies in other asset classes, see Lima et al. (2020) and reference therein, and are in favour of the efficient market hypothesis of Fama (1970) for short-term forecasting horizons.

The comparison with BM2 allows us to refine our analysis and determine, among other things, if the out-performance of the alternative models with respect to the BM1 is due to lagged values of the dependent variables or to the additional indicators obtained from the data described in Section 4. A casual inspection of Figures 6-9 highlights the following points. First, except for some cases, most of the variables are informative for future retail price variations. But some classes of variables are more important than others. Models based on Michigan consumer surveys, order inventories, interest rates and exchange rates data deliver the best performances, especially for horizons above one year. While our models have no theoretical economic underpinning, some of these results are in line with intuitions. For example, consumers, in performing their projections, implicitly account for how their demand for goods might evolve in the near future, since future economic conditions are linked to the evolution of their wealth. Given the non-negligible price elasticity of wine and beer to income, see Samuelson et al. (2021), this would vary the demand and hence the equilibrium prices. Wholesale prices are predicted by similar variables, but in a milder way, i.e. the forecasting power is mainly concentrated for horizons of two years. Second, the future fine wine prices (both LIVX50 and LIVX100) can be foretasted for horizons of two years, mainly with monetary variables (M6), financial variables - risk factors on mature, developed markets, M19, and liquidity conditions, M21- and international variables linked to industrial production and trade conditions (M15 to M19). The finding that fine wine prices are linked to internationally-related economic and financial variables might be due to the fact that fine wines might be seen as alternative investments, and as such, their prices might be linked to general financial markets patterns such as (i) fundamental risks in the stock market, proxied by the Fama-French factors, (ii) exchange rates dynamics, and (iii) global economic conditions.

As a final analysis, we asses if, by optimally combining models, one can improve the forecasting performances. In the second panel group (Figures 10-17), the horizontal axes report the combined forecasts, considering the groups of explanatory variables *Gs*. Similarly to the stand-alone model analysis, the ten *Gs* combined forecasts are compared to the two benchmark models (BM1 and BM2). In general, the optimal combination leads to a generalised improvement, i.e. across dependent variables, in the forecasting results, but with some important heterogeneity. When forecasting wholesale and retail prices, the optimal combination increases the forecasting power for most of the explanatory variables at all forecasting horizons. However, for forecasting fine wine prices, the improvement in the forecast is especially elevated for short-term horizons.

Overall, our results indicate that forecasting retail, wholesale and fine wine prices are possible, but it is crucial to carefully select the explanatory variables.

6 Conclusions

Forecasting wine prices is relevant for market operators to define and plan consumption, selling and investment strategies. This study has examined the performance of different forecasting models of fine wine, alcohol retail and wholesale price changes and compared them to a set of benchmarks for the period 1992-2022 using monthly data.

We show that, depending on the forecasting horizons, it is possible to forecast with a certain accuracy fine wine and alcohol price. Fine wines prices can be forecasted for horizons of two years, while wholesale and retail prices can be forecasted also for shorter horizons. The findings indicate that international aggregate economic indicators and equity risk factors of mature markets are key for predicting fine wine prices while local (US) factors included in McCracken and Ng (2020) and consumer surveys are crucial for forecasting retail and wholesale prices. Our results are particular relevant for financial investors seeking extra returns in alternative assets and to any economic actor business activity is linked to alcohol and wine prices.

Our analysis indicate several interesting new avenue for research. First, it would be interesting to assess whether weather conditions, such as global temperature anomalies and the El-Nino cycles, can predict long term (fine) wine prices and wine quality. Second, our forecasting setting is mainly static and based on direct techniques. It would be interesting to know if other, more refines, settings might lead to better outcomes. We pin these ideas and plan to develop them further in the near future.

References

Ameur, H. B., Z. Ftiti, and E. Le Fur (2022). What can we learn from the analysis of the fine wines market efficiency? *International Journal of Finance & Economics n/a*(n/a), 1–16.

Ashenfelter, O. (2008). Predicting the quality and prices of Bordeaux wine. *The Economic Journal 118*(529), F174–F184.

Ashenfelter, O. and K. Storchmann (2010). Measuring the economic effect of global warming on viticulture using auction, retail, and wholesale prices. *Review of Industrial Organization 37*(1), 51–64.

Baker, S. R., N. Bloom, S. J. Davis, and K. J. Kost (2019, March). Policy news and stock market volatility. Working Paper 25720, National Bureau of Economic Research.

Bazen, S. and J. Cardebat (2021). Forecasting Bordeaux wine prices using state space methods. *Applied Economics 40*(47).

Bazen, S. and J. Cardebat (2022). Why have Bordeaux wine prices become so difficult to forecast? *Economics Bulletin 42*(1).

Bouri, E. I. and D. Roubaud (2016). Fine wines and stocks from the perspective of uk investors: Hedge or safe haven? *Journal of Wine Economics 11*(2), 233–248.

Burton, B. J. and J. P. Jacobsen (1999). Measuring returns on investments in collectibles. *The Journal of Economic Perspectives 13*(4), 193–212.

Cardebat, J. M., B. Faye, E. L. Fur, and K. Storchmann (2017). The law of one price? price dispersion on the auction market for fine wine. *Journal of Wine Economics 12*(3).

Cardebat, J.-M. and L. Jiao (2018). The long-term financial drivers of fine wine prices: The role of emerging markets. *The Quarterly Review of Economics and Finance 67*, 347–361.

Cross, R., A. J. Plantinga, and R. N. Stavins (2011, May). What is the value of terroir? *American Economic Review 101*(3), 152–56.

Diebold, F. X. and R. S. Mariano (2002). Comparing predictive accuracy. *Journal of Business & Economic Statistics 20*(1), 134–144.

Dimson, E., P. L. Rousseau, and C. Spaenjers (2015). The price of wine. *Journal of Financial Economics 118*(2), 431–449.

Elliott, G. and A. Timmermann (2016). Forecasting in economics and finance. *Annual Review of Economics 8*(1), 81–110.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance 25*(2), 383–417.

Fama, E. F. and K. R. French (1987). Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *The Journal of Business 60*(1), 55–73.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Faust, J. and J. H. Wright (2013). Chapter 1 - forecasting inflation. In G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2 of *Handbook of Economic Forecasting*, pp. 2–56. Elsevier.

Faye, B., E. L. F. E., and S. Prat (2015). Dynamics of fine wine and asset prices: evidence from short-and long-run co-movements. *Applied Economics 47*(29).

Fogarty, J. J. (2006, December). The return to Australian fine wine. *European Review of Agricultural Economics 33*(4), 542–561.

Fogarty, J. J. and R. Sadler (2014). To save or savor: A review of approaches for measuring wine as an investment. *Journal of Wine Economics 9*(3).

Gergaud, O. and V. Ginsburgh (2008). Natural endowments, production technologies and the quality of wines in Bordeaux. does terroir matter? *The Economic Journal 118*(529), F142–F157.

Gergaud, O., A. J. Plantinga, and A. Ringeval-Deluze (2017). Anchored in the past: Persistent price effects of obsolete vineyard ratings in France. *Journal of Economic Behavior & Organization 133*(C), 39–51.

Haeck, C., G. Meloni, and J. Swinnen (2019). The value of terroir: A historical analysis of the Bordeaux and champagne geographical indications. *Applied Economic Perspectives and Policy 41*(4), 598–619.

Le Fur, E., H. Ben Ameur, and B. Faye (2016). Time-varying risk premiums in the framework of wine investment. *Journal of Wine Economics 11*(3), 355–378.

Lima, L. R., F. Meng, and L. Godeiro (2020). Quantile forecasting with mixed-frequency data. *International Journal of Forecasting 36*(3), 1149–1162.

Martínez-García, E., V. Grossman, and A. Mack (2015). A contribution to the chronology of turning points in global economic activity (1980–2012). *Journal of Macroeconomics 46*, 170–185.

Masset, P. and C. Henderson (2010). Wine as an alternative asset class. *Journal of Wine Economics 5*(1), 87–118.

Masset, P. and J.-P. Weisskopf (2022). At what price should Bordeaux wines be released? *Economic Inquiry 60*(1), 392–412.

McCracken, M. W. and S. Ng (2020). Fred-qd: A quarterly database for macroeconomic research. *Federal Reserve Bank of St. Louis Working Paper* (005).

Mundi, H. S. and D. Kumar (2023). The potential of alternative investments as an asset class: a thematic and bibliometric review. *Qualitative Research in Financial Markets 15*(1), 119–141.

Nahmer, T. (2020). Diversification benefit of actual investing in fine wine. *The Journal of Alternative Investments 22*(4), 59–74.

Oczkowski, E. (2001). Hedonic wine price functions and measurement error. *Economic Record 77*(239), 374–382.

Oleksy, P., M. Czupryna, and M. Jakubczyk (2021). On fine wine pricing across different trading venues. *Journal of Wine Economics 16*(2), 189–209.

Paroissien, E. (2020). Forecasting bulk prices of Bordeaux wines using leading indicators. *International Journal of Forecasting 36*(2), 292–309.

Pástor, L. and R. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy 111*(3), 642–685.

Samuelson, W. F., S. G. Marks, and J. L. Zagorsky (2021). *Managerial Economics, 9th edition*. Wiley.

Yeo, M., T. Fletcher, and J. Shawe-Taylor (2015). Machine learning in fine wine price prediction. *Journal of Wine Economics 10*(2), 151–172.

A Commodities

- Food and Beverage Price Index, 2016 = 100, includes Food and Beverage Price Indices
- Food Price Index, 2016 = 100, includes Cereal, Vegetable Oils, Meat, Seafood, Sugar, and Other Food (Apple (non-citrus fruit), Bananas, Chana (legumes), Fishmeal, Groundnuts, Milk (dairy), Tomato (veg)) Price Indices
- Beverage Price Index, 2016 = 100, includes Coffee, Tea, and Cocoa
- Industrial Inputs Price Index, 2016 = 100, includes Agricultural Raw Materials and Base Metals Price Indices
- Agriculture Price Index, 2016 = 100, includes Food and Beverages and Agriculture Raw Materials Price Indices
- Agricultural Raw Materials Index, 2016 = 100, includes Timber, Cotton, Wool, Rubber, and Hides Price Indices
- All Metals Index, 2016 = 100: includes Metal Price Index (Base Metals) and Precious

Metals Index

- Base Metals Price Index, 2016 = 100, includes Aluminum, Cobalt, Copper, Iron Ore, Lead, Molybdenum, Nickel, Tin, Uranium and Zinc Price Indices

- Precious Metals Price Index, 2016 = 100, includes Gold, Silver, Palladium and Platinum Price Indices

- All Metals EX GOLD Index, 2016 = 100: includes Metal Price Index (Base Metals) and ONLY Silver, Palladium, Platinum

- Fuel (Energy) Index, 2016 = 100, includes Crude oil (petroleum), Natural Gas, Coal Price and Propane Indices

- Crude Oil (petroleum), Price index, 2016 = 100, simple average of three spot prices; Dated Brent, West Texas Intermediate, and the Dubai Fateh

- Natural Gas Price Index, 2016 = 100, includes European, Japanese, and American Natural Gas Price Indices

- Coal Price Index, 2016 = 100, includes Australian and South African Coal

B Tables

Table 1: Model Description

Model	Explanatory Variables
M1	Output income
M2	Labor market
M3	Consumption
M4	Orders inventories
M5	Money credit
M6	Interest rates and exchanges
M7	Prices
M8	Stock Market
M9	EMV
M10	Michigan all
M11	Michigan icc
M12	Michigan ics
M13	Michigan ice
M14	cpi
M15	ip
M16	imports
M17	ppi
M18	ner
M19	Developed 5 Factors
M20	Emerging 5 Factors
M21	Liquidity
M22	Commodities

Table 2: Combined forecasts: Groups of Explanatory Variables

Group	Description
G1	Output income+Labor market+Consumption+Orders inventories+ Money credit+Interest rates and exchanges+Prices+Stock Market
G2	Michigan icc+Michigan ics+Michigan ice
G3	cpi+ip+imports+ppi+ner
G4	Developed 5 Factors+Emerging 5 Factors
G5	Output income+Labor market+Consumption+Orders inventories
G6	Money credit+Interest rates and exchanges+Prices+Stock Market
G7	Interest rates and exchanges+Prices+Stock Market+ Developed 5 Factors+Emerging 5 Factors
G8	Developed 5 Factors+Emerging 5 Factors+Commodities
G9	All Macro+Commodities

Table 3: Percentage of cases for which the score ratio between the RMSE of a model (or combination of models) and the one of the benchmark model BM2 is less than one.

	Wholesale				Retail				LIVX100				LIVX50			
	Lasso	EN	Ridge	PCA	Lasso	EN	Ridge	PCA	Lasso	EN	Ridge	PCA	Lasso	EN	Ridge	PCA
$h = 1$																
$p = 1, q = 1$	16.13	12.90	3.23	64.52	70.97	64.52	58.06	87.10	0.00	3.23	0.00	16.13	16.13	12.90	6.45	12.90
$p = 1, q = 3$	12.90	9.68	0.00	54.84	38.71	41.94	16.13	74.19	0.00	3.23	0.00	6.45	9.68	3.23	0.00	3.23
$p = 3, q = 1$	3.23	0.00	0.00	16.13	38.71	41.94	6.45	19.35	0.00	3.23	0.00	6.45	19.35	16.13	6.45	0.00
$p = 3, q = 3$	3.23	0.00	0.00	16.13	12.90	19.35	0.00	6.45	0.00	0.00	0.00	3.23	16.13	6.45	0.00	3.23
$h = 6$																
$p = 1, q = 1$	3.23	0.00	0.00	0.00	90.32	90.32	83.87	83.87	16.13	12.90	19.35	32.26	0.00	0.00	0.00	9.68
$p = 1, q = 3$	0.00	0.00	0.00	9.68	87.10	87.10	83.87	77.42	16.13	16.13	12.90	19.35	0.00	0.00	0.00	3.23
$p = 3, q = 1$	12.90	16.13	16.13	22.58	64.52	67.74	61.29	12.90	19.35	16.13	19.35	29.03	0.00	0.00	0.00	6.45
$p = 3, q = 3$	16.13	3.23	0.00	29.03	83.87	83.87	70.97	19.35	16.13	12.90	16.13	19.35	0.00	0.00	0.00	3.23
$h = 12$																
$p = 1, q = 1$	19.35	25.81	25.81	29.03	67.74	77.42	77.42	83.87	3.23	3.23	3.23	9.68	0.00	0.00	0.00	0.00
$p = 1, q = 3$	25.81	22.58	22.58	32.26	74.19	77.42	74.19	77.42	3.23	3.23	3.23	3.23	0.00	0.00	3.23	0.00
$p = 3, q = 1$	22.58	32.26	29.03	32.26	32.26	32.26	48.39	35.48	3.23	3.23	3.23	9.68	0.00	0.00	0.00	3.23
$p = 3, q = 3$	29.03	22.58	22.58	32.26	41.94	51.61	54.84	29.03	3.23	3.23	3.23	3.23	0.00	0.00	3.23	6.45
$h = 18$																
$p = 1, q = 1$	38.71	35.48	35.48	38.71	64.52	70.97	61.29	74.19	19.35	19.35	16.13	35.48	22.58	12.90	12.90	25.81
$p = 1, q = 3$	29.03	29.03	32.26	35.48	61.29	64.52	64.52	74.19	16.13	12.90	19.35	25.81	19.35	19.35	12.90	19.35
$p = 3, q = 1$	38.71	41.94	38.71	51.61	61.29	54.84	61.29	64.52	19.35	16.13	16.13	29.03	16.13	16.13	9.68	19.35
$p = 3, q = 3$	32.26	32.26	32.26	45.16	58.06	58.06	58.06	61.29	19.35	19.35	19.35	25.81	19.35	19.35	9.68	12.90
$h = 24$																
$p = 1, q = 1$	58.06	58.06	58.06	64.52	74.19	74.19	64.52	80.65	32.26	32.26	32.26	29.03	29.03	29.03	35.48	38.71
$p = 1, q = 3$	61.29	58.06	58.06	61.29	64.52	64.52	58.06	67.74	19.35	19.35	19.35	19.35	29.03	29.03	32.26	29.03
$p = 3, q = 1$	48.39	48.39	48.39	45.16	58.06	61.29	61.29	64.52	19.35	19.35	16.13	19.35	29.03	32.26	29.03	22.58
$p = 3, q = 3$	51.61	54.84	54.84	58.06	64.52	61.29	61.29	64.52	19.35	19.35	19.35	12.90	29.03	32.26	29.03	25.81

Note: For each forecasting horizon, estimation method, choice of lags in eq. (2), and wine index, the table reports the percentage of models (or combination of models) that have a RMSE smaller than benchmark model BM2.

C Figures

Figure 2: Heatmap for the Score Ratios. Index: Liv-ex Fine Wine 50. Method: lasso, Benchmark: BM1.

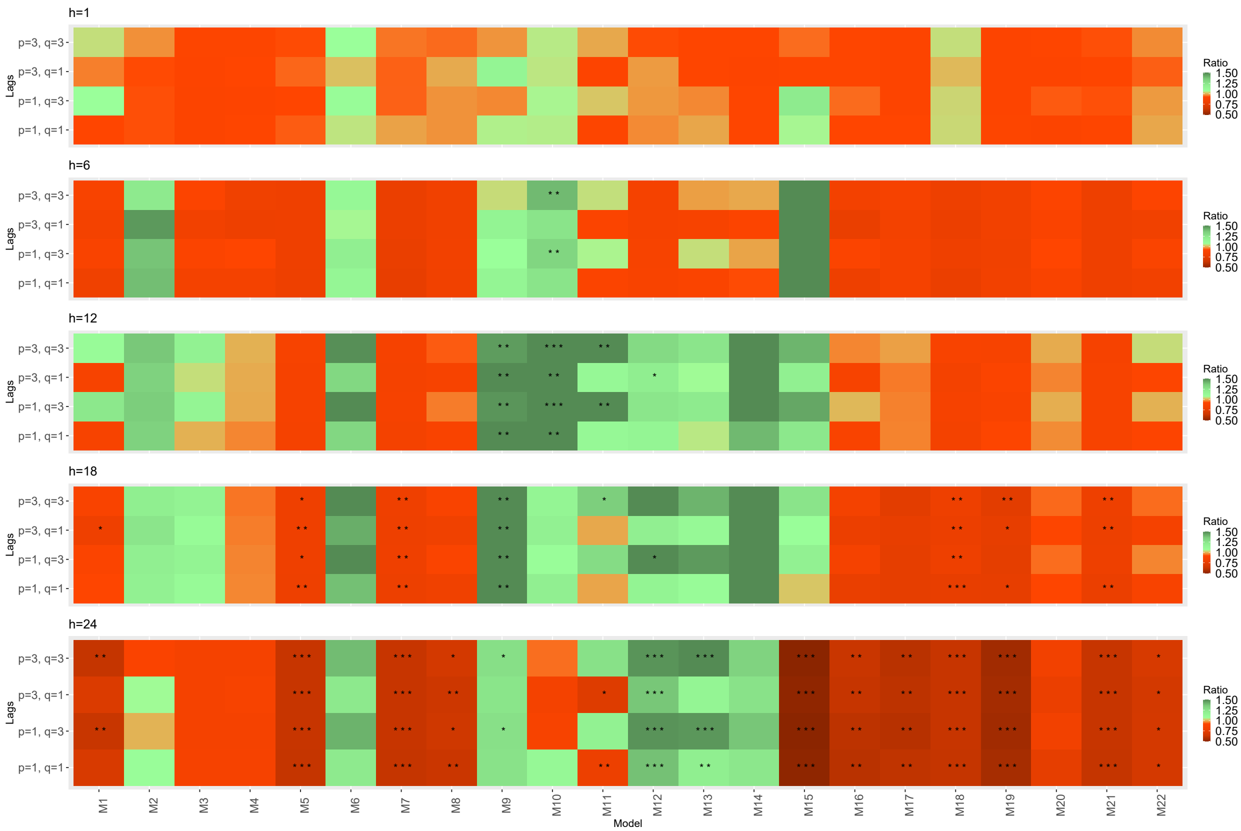

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 3: Heatmap for the Score Ratios. Index: Liv-ex Fine Wine 100. Method: lasso, Benchmark: BM1.

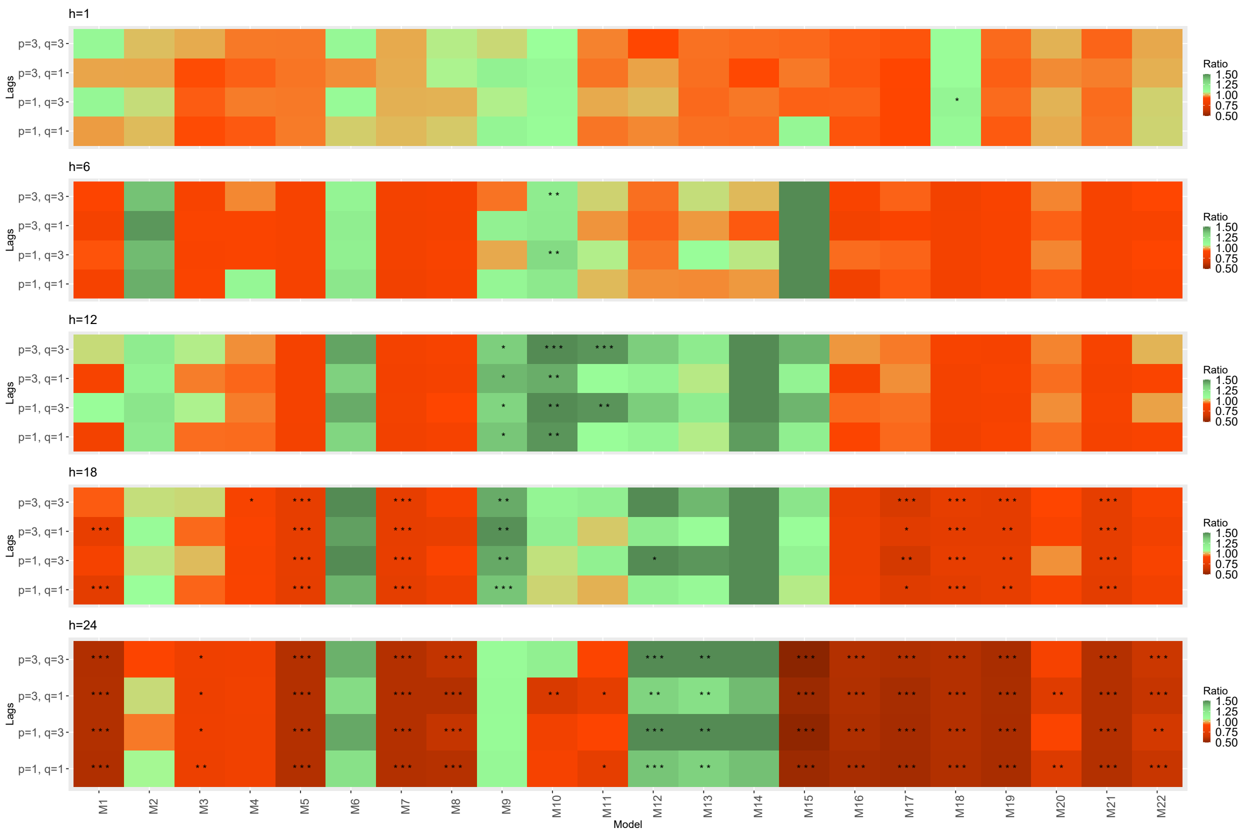

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 4: Heatmap for the Score Ratios. Index: Retail. Method: lasso, Benchmark: BM1.

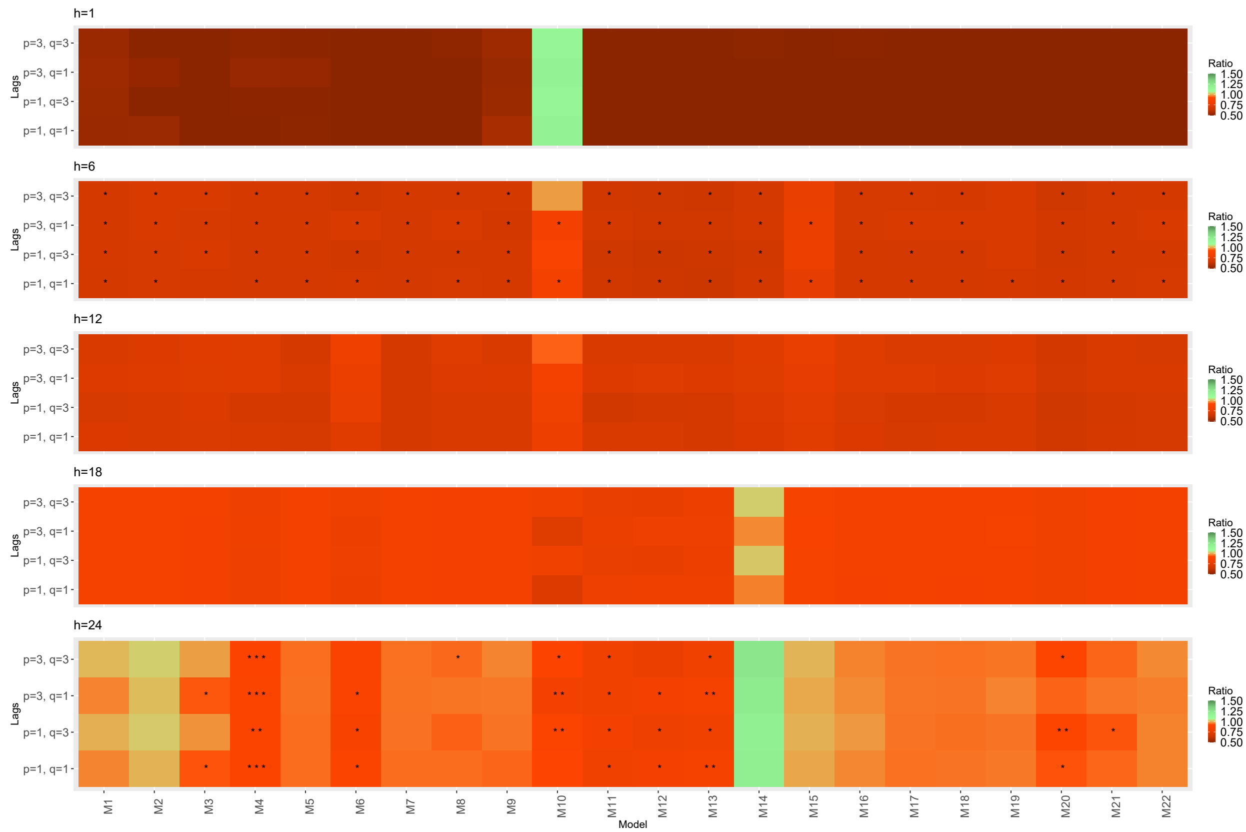

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 5: Heatmap for the Score Ratios. Index: Wholesale. Method: lasso, Benchmark: BM1.

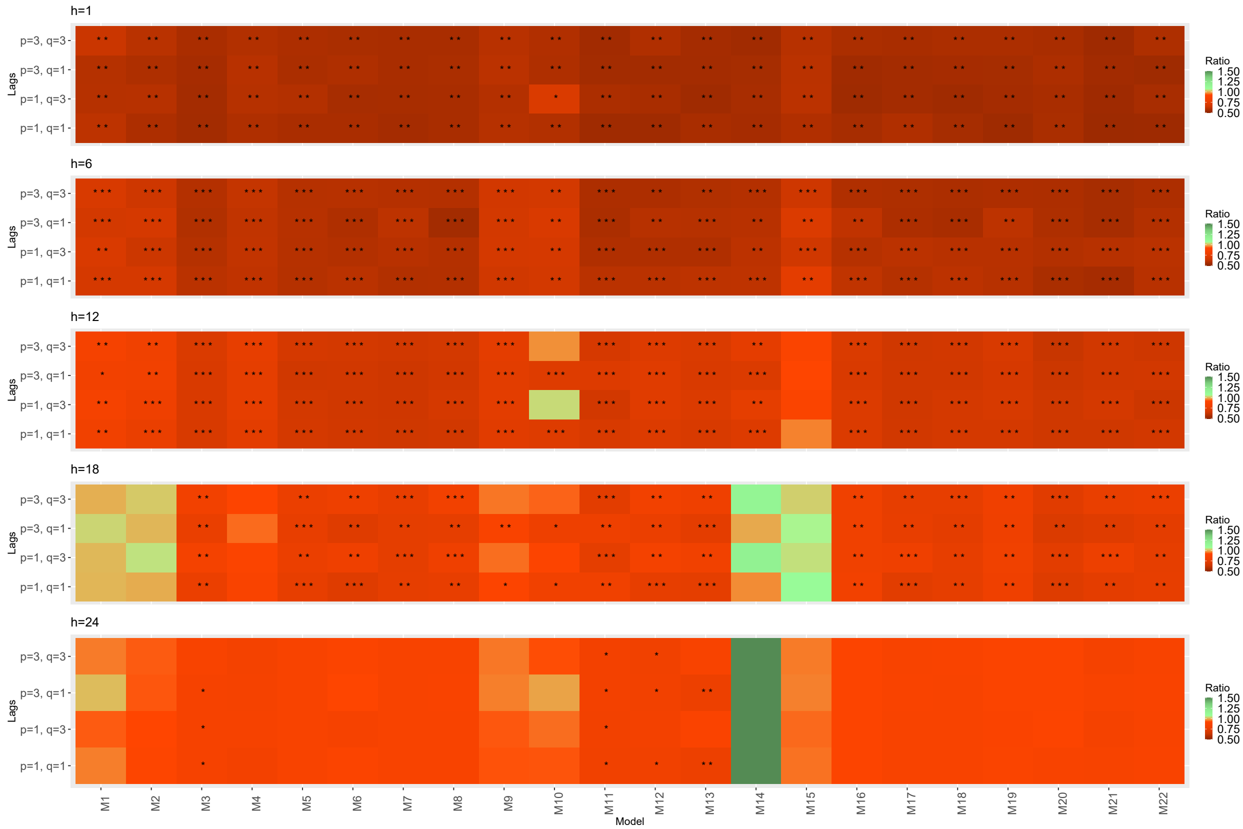

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 6: Heatmap for the Score Ratios. Index: Liv-ex Fine Wine 50. Method: lasso, Benchmark: BM2.

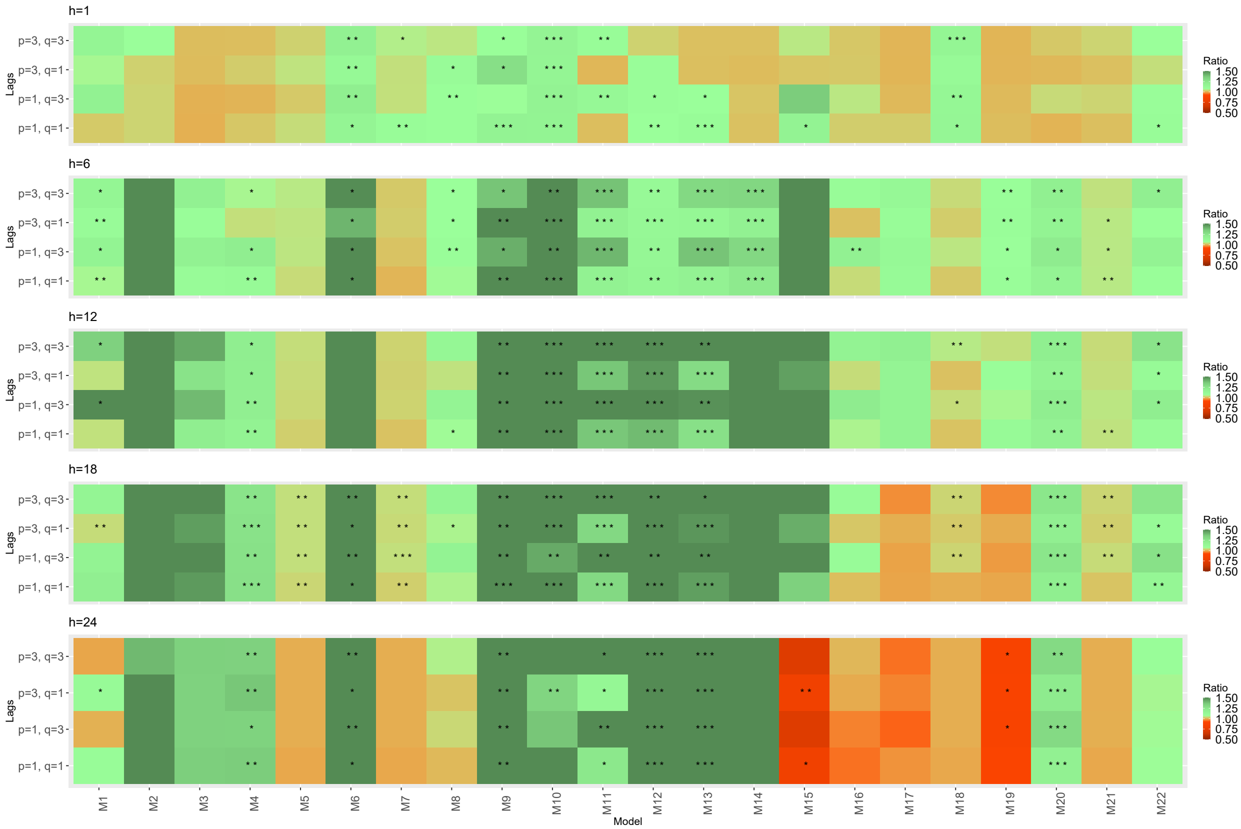

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 7: Heatmap for the Score Ratios. Index: Liv-ex Fine Wine 100. Method: lasso, Benchmark: BM2.

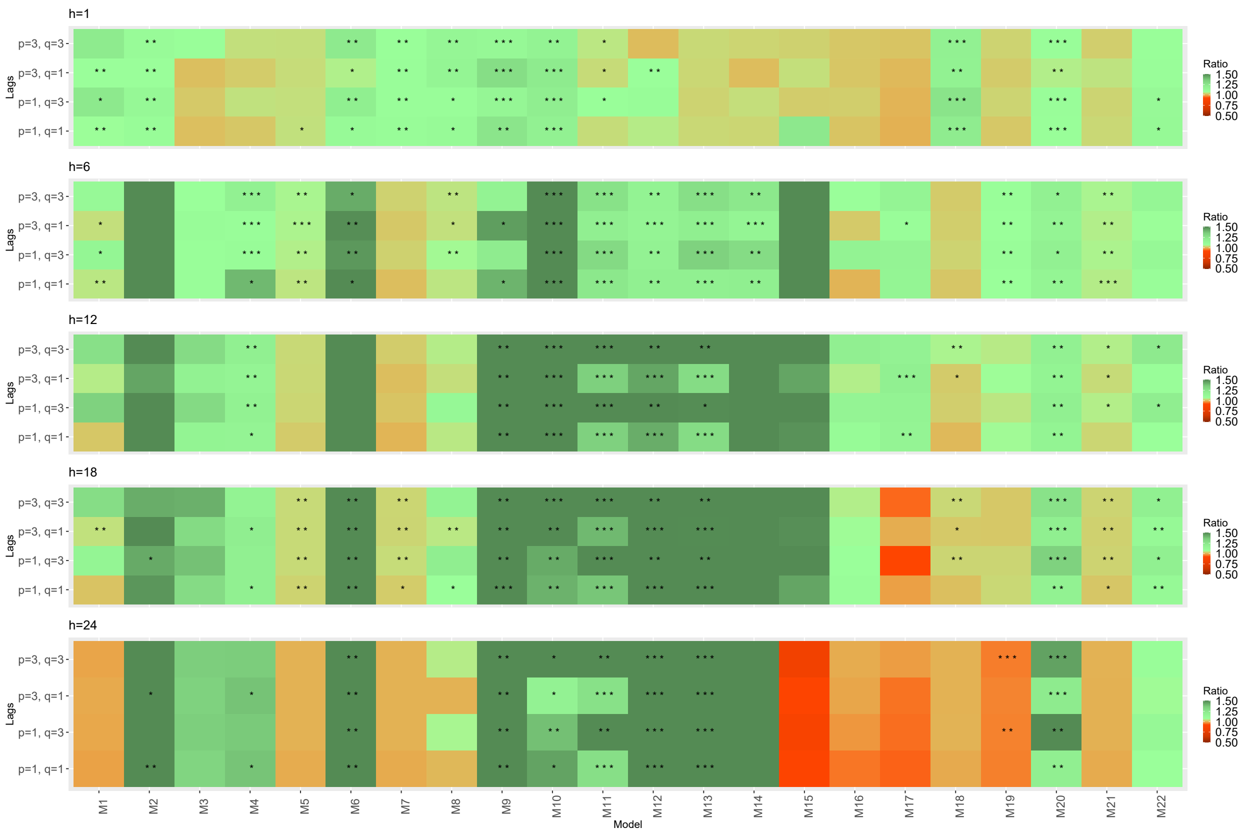

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 8: Heatmap for the Score Ratios. Index: Retail. Method: lasso, Benchmark: BM2.

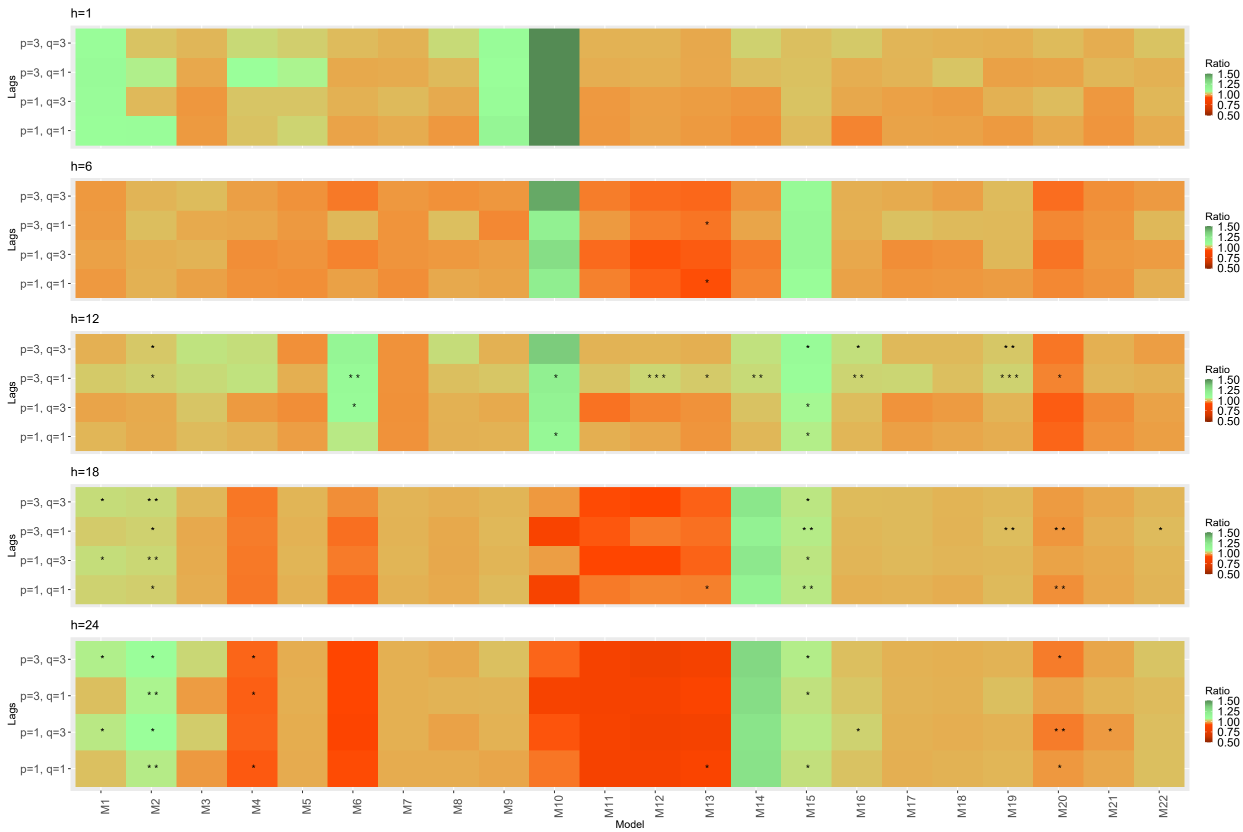

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 9: Heatmap for the Score Ratios. Index: Wholesale. Method: lasso, Benchmark: BM2.

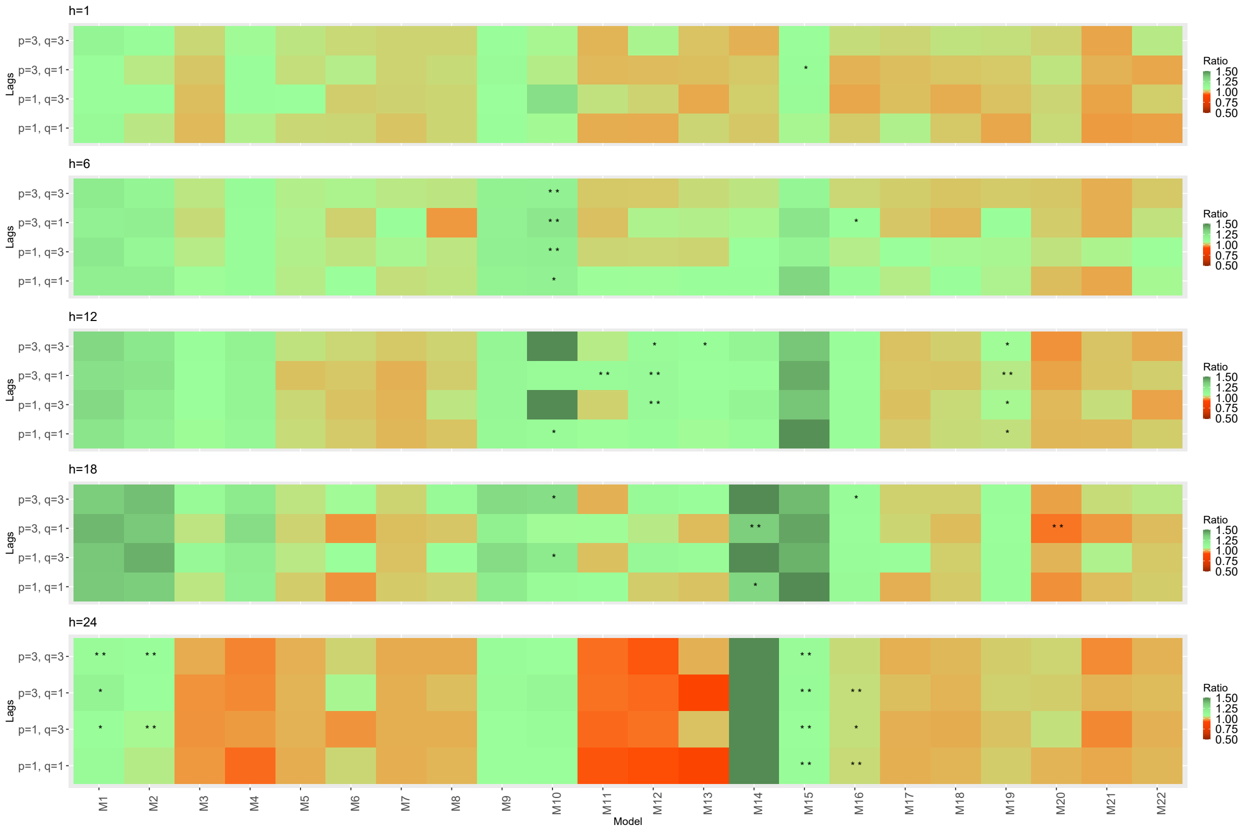

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 10: Heatmap for the Score Ratios (Combined Forecasts). Index: Liv-ex Fine Wine 50. Method: lasso, Benchmark: BM1.

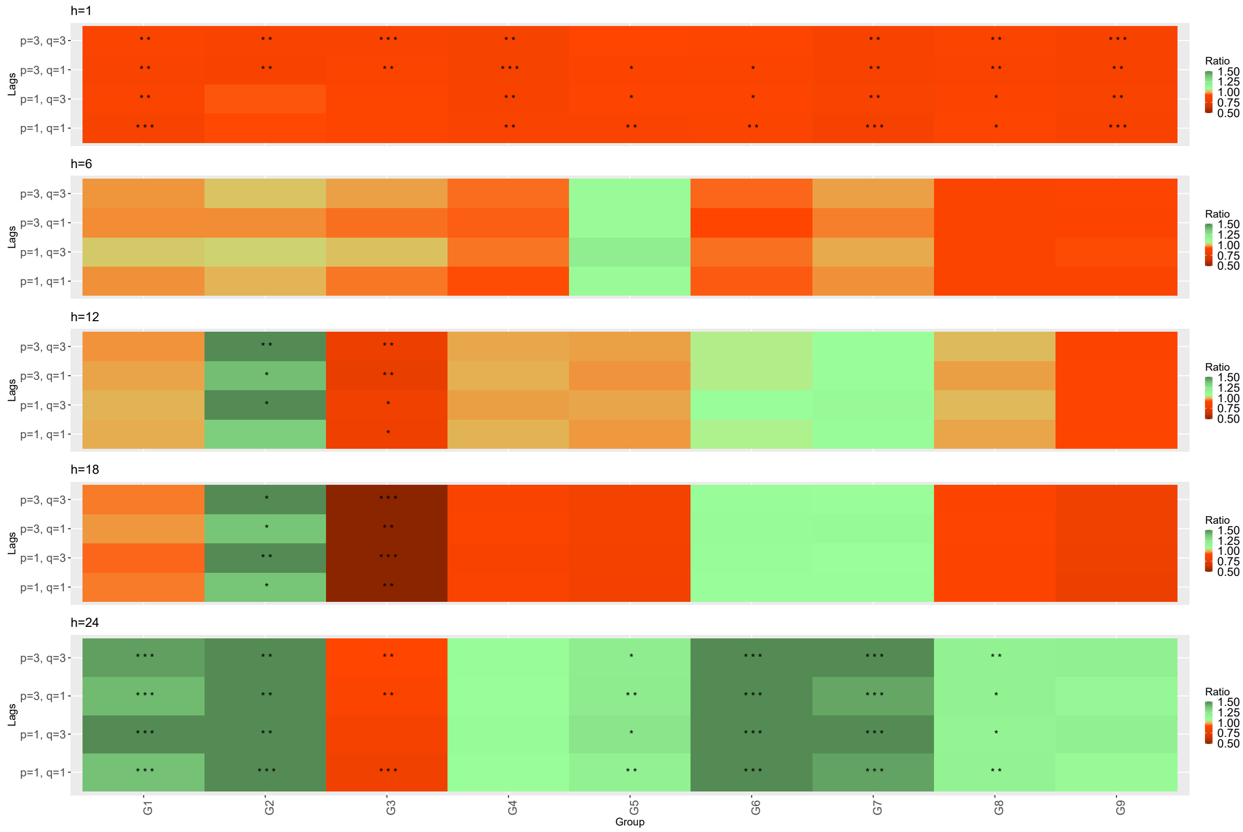

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 11: Heatmap for the Score Ratios (Combined Forecasts). Index: Liv-ex Fine Wine 100. Method: lasso, Benchmark: BM1.

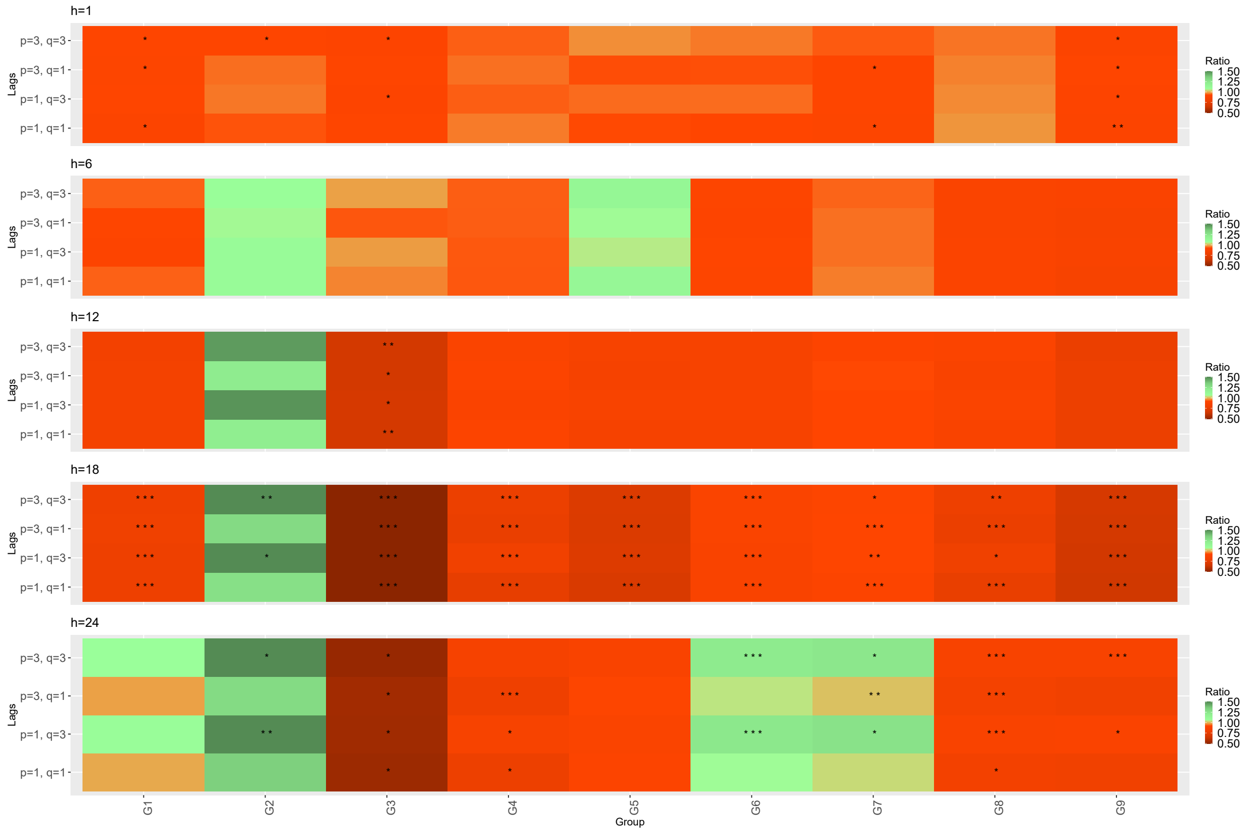

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 12: Heatmap for the Score Ratios (Combined Forecasts). Index: Retail. Method: lasso, Benchmark: BM1.

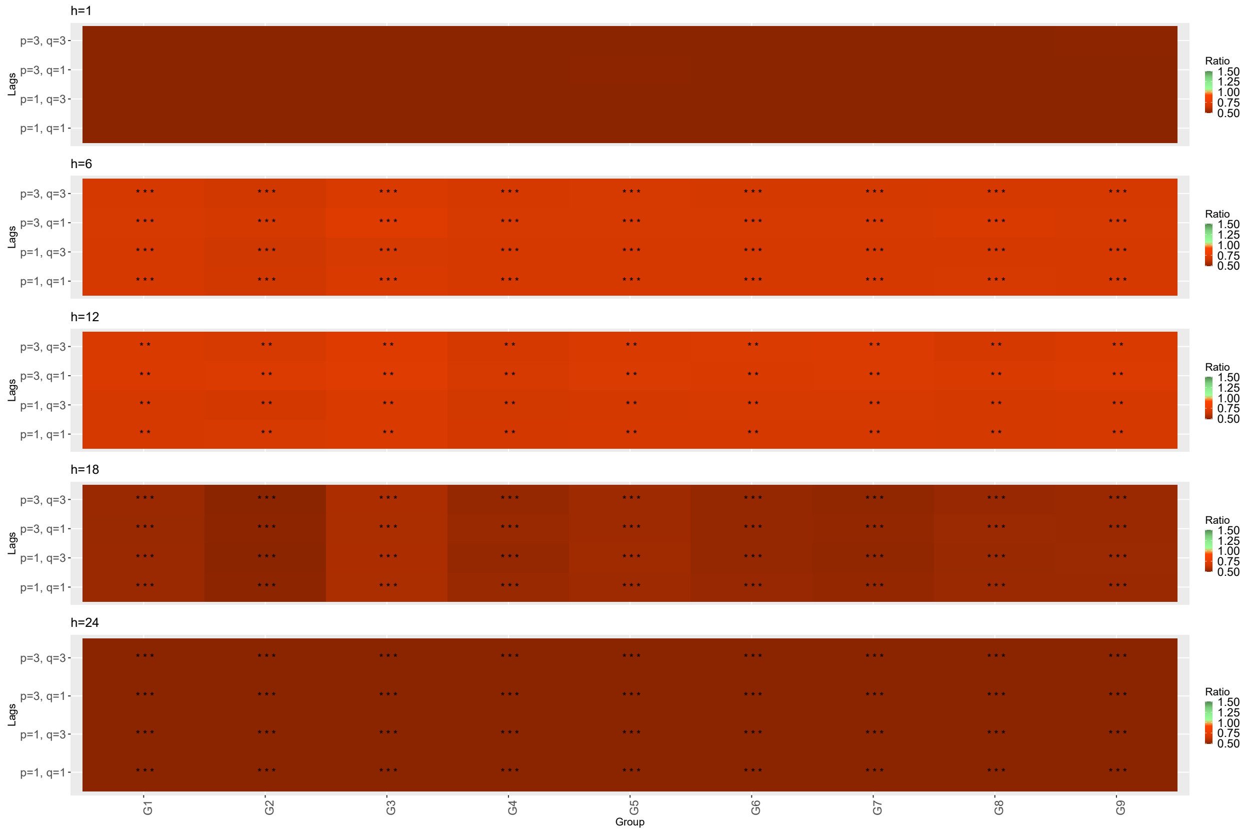

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 13: Heatmap for the Score Ratios (Combined Forecasts). Index: Wholesale. Method: lasso, Benchmark: BM1.

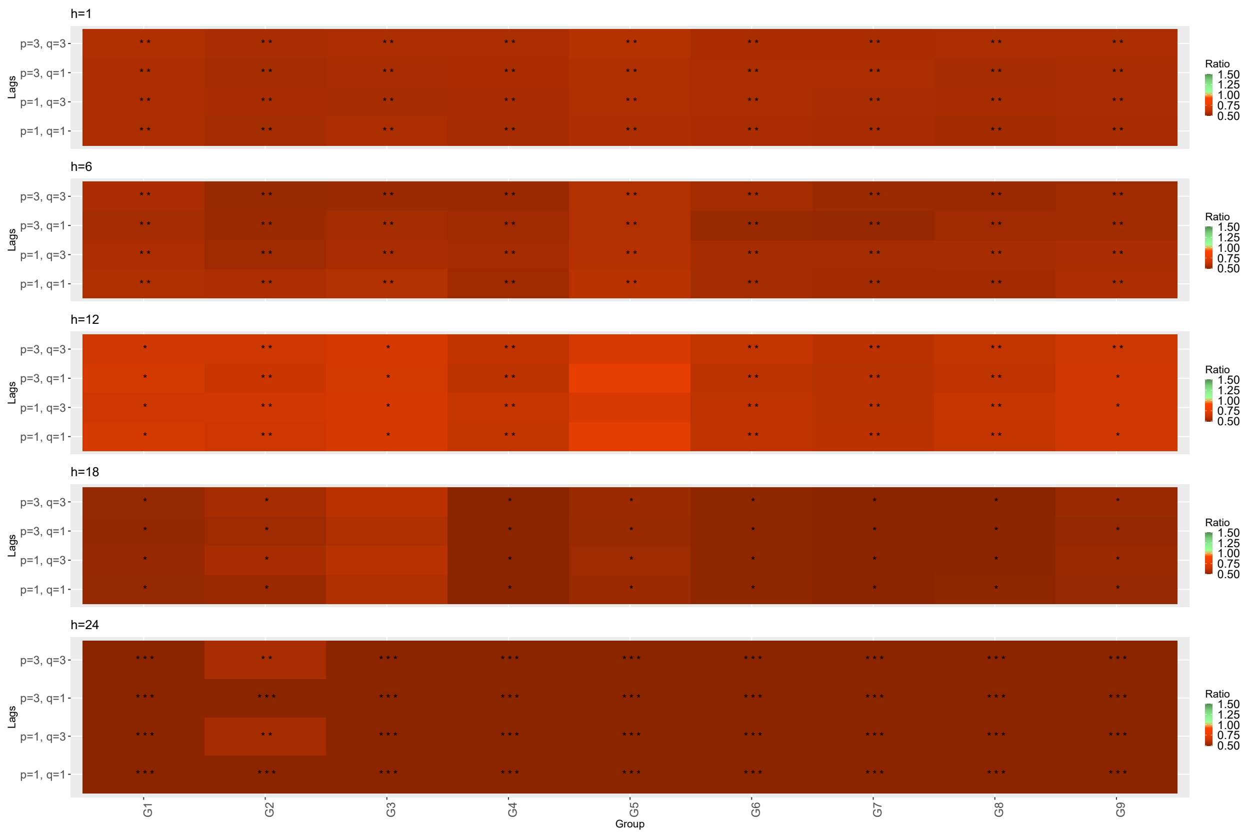

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM1 (Random walk in growth). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 14: Heatmap for the Score Ratios (Combined Forecasts). Index: Liv-ex Fine Wine 50. Method: lasso, Benchmark: BM2.

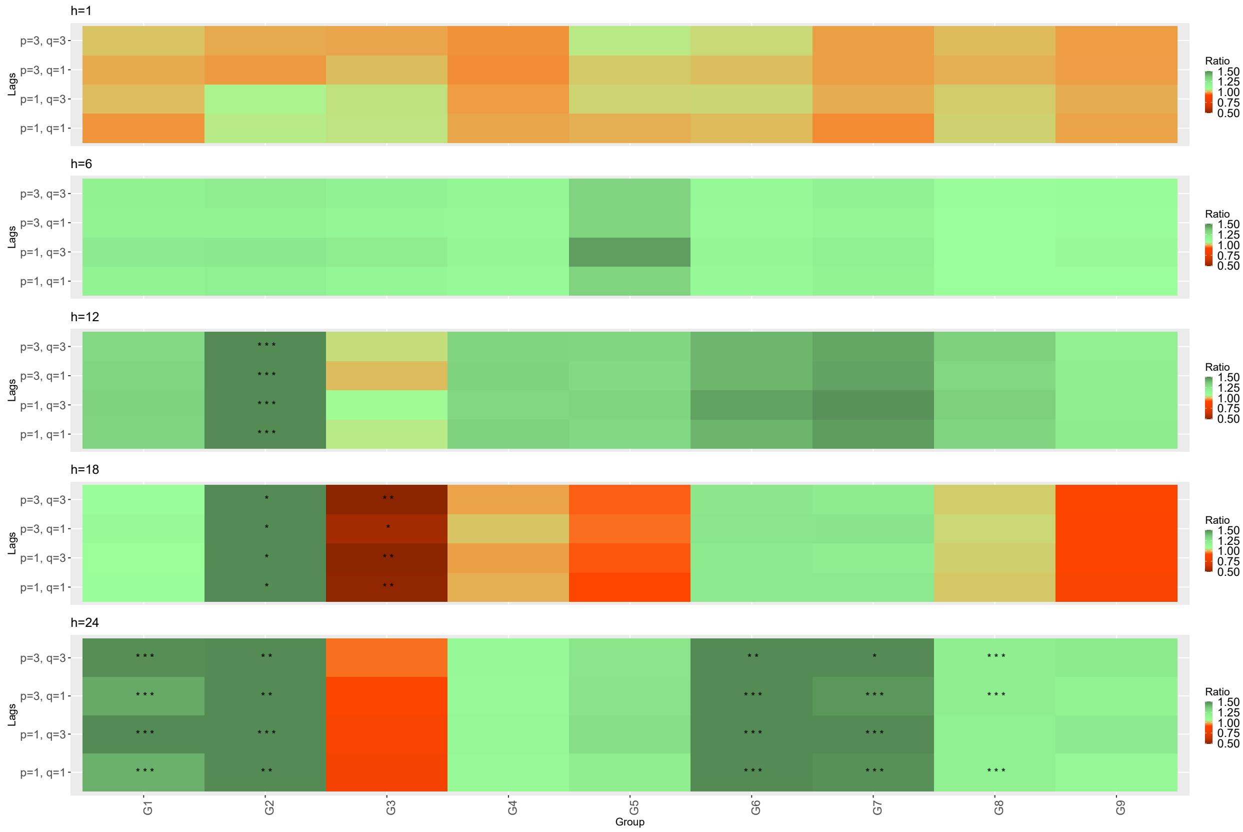

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 15: Heatmap for the Score Ratios (Combined Forecasts). Index: Liv-ex Fine Wine 100. Method: lasso, Benchmark: BM2.

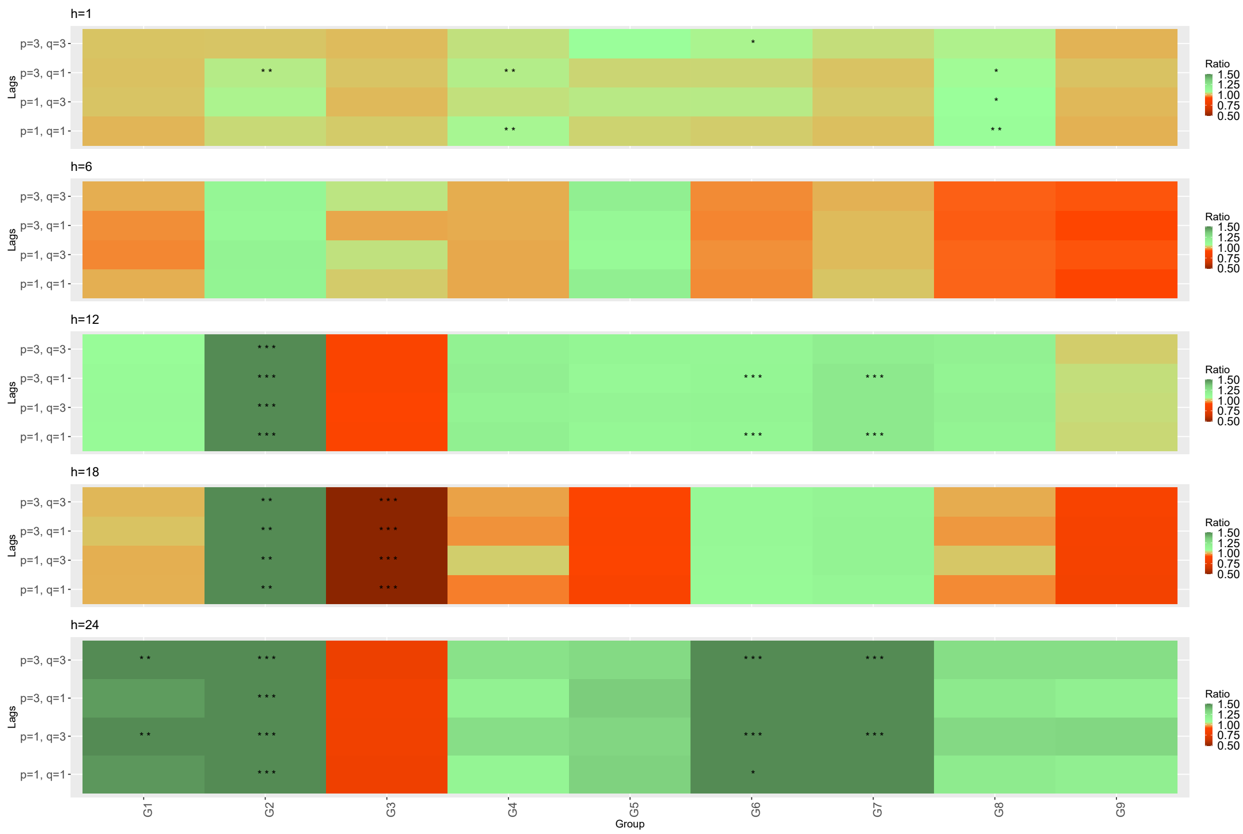

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 16: Heatmap for the Score Ratios (Combined Forecasts). Index: Retail. Method: lasso, Benchmark: BM2.

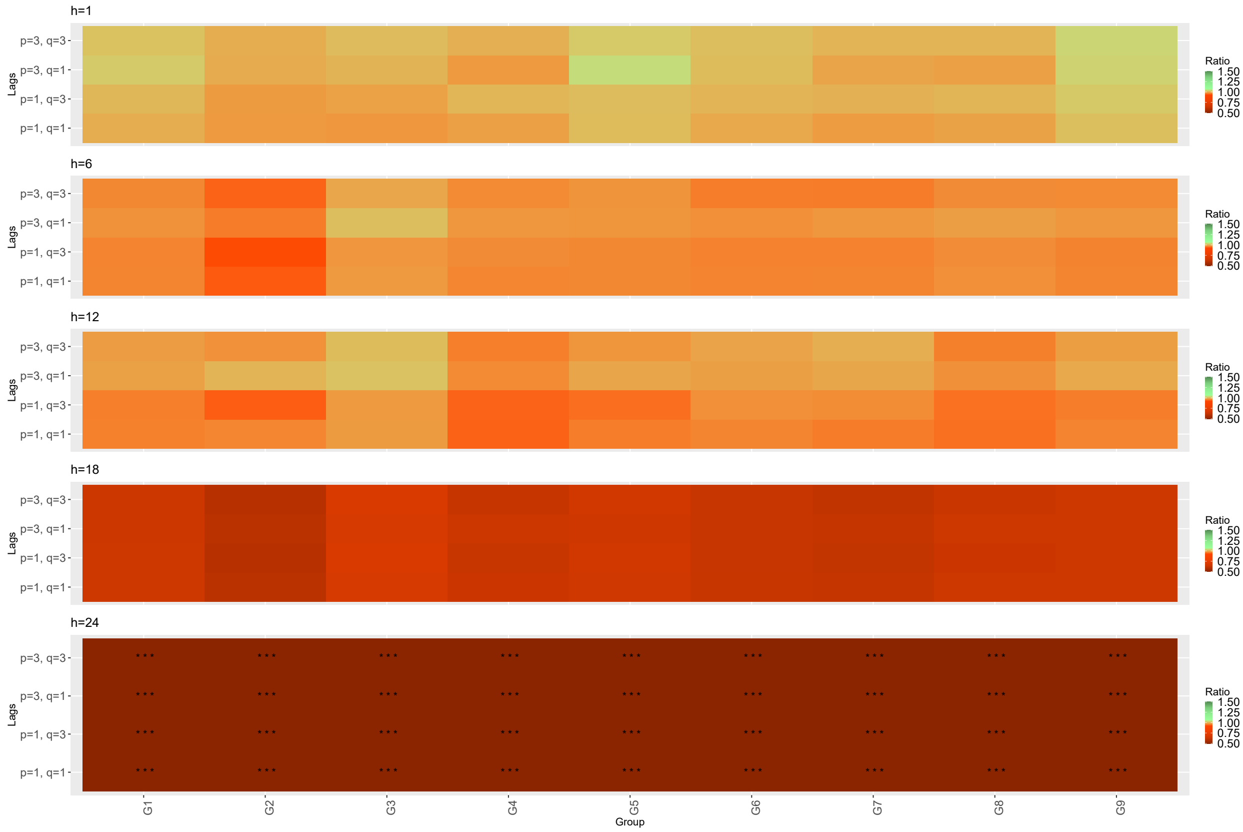

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.

Figure 17: Heatmap for the Score Ratios (Combined Forecasts). Index: Wholesale. Method: lasso, Benchmark: BM2.

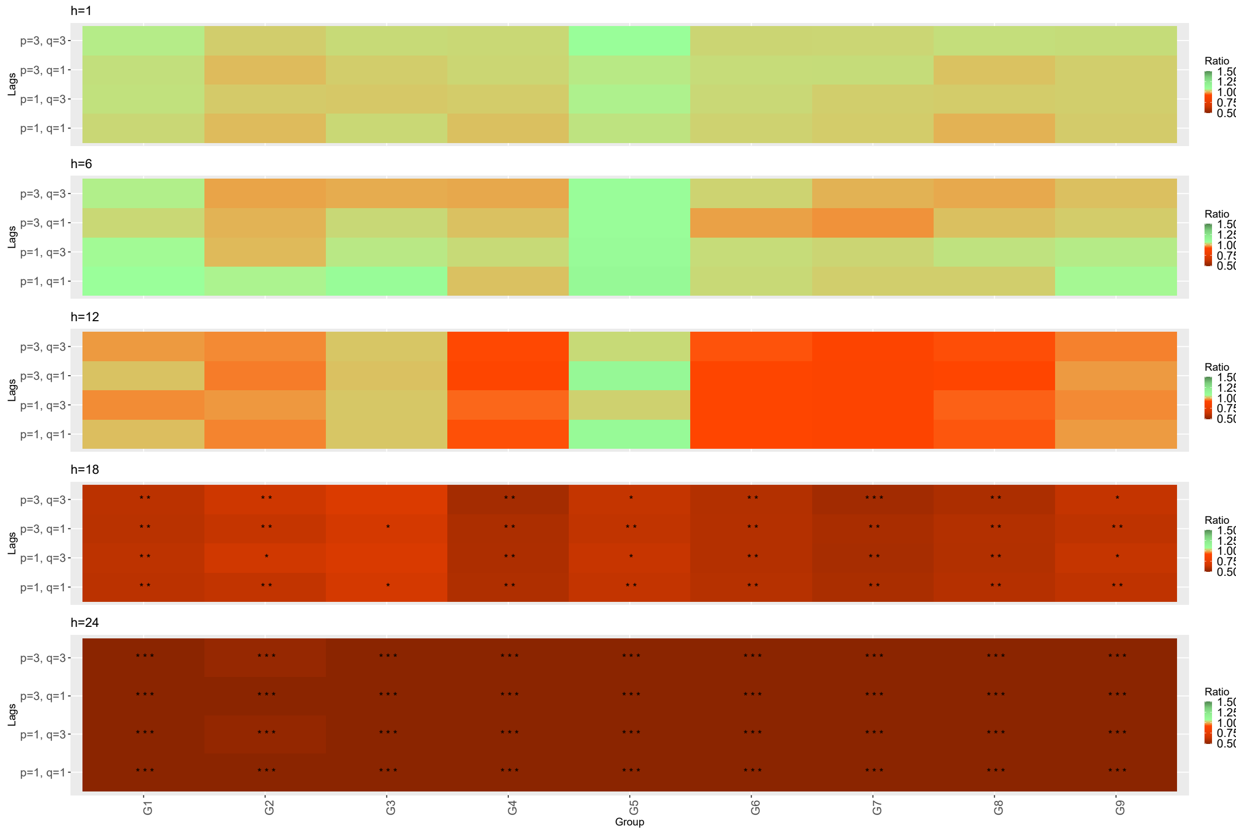

Note: SR is the ratio between the RMSE of a given model and the one of the benchmark model BM2 (Distributed lag model with the last three lagged month on month changes in wine prices as control variables). The stars in each cell are related to the p-values for the null hypothesis of equal predictability (test of Diebold and Mariano, 2002). ***, **, and *, denote significance at the 1%, 5% and 10% levels, respectively.