Journal of Wine Economics, Volume 14, Number 2, 2019, Pages 208–213
doi:10.1017/jwe.2019.15

A Note on Forecasting Alcohol Demand

Derby Voon [a] **and James Fogarty** [b]

Abstract

A recent study in the *Journal of Wine Economics* presented forecasts of future alcohol consumption derived using the ARIMA (Box–Jenkins) method. Alcohol consumption forecasts can be developed using many different methodologies. In this Note we highlight the value of using multiple methods to develop alcohol consumption forecasts, and demonstrate the capability of the **R** software platform as a general forecasting tool. (JEL Classifications: D12, C53)

Keywords: alcohol consumption, forecasting.

I. Introduction

Convergence in global alcohol consumption patterns is an active research area (Smith and Mitry, 2007; Aizenman and Brooks, 2008; Colen and Swinnen, 2016; Holmes and Anderson, 2017), as is how to measure convergence in alcohol consumption patterns (Mills, 2018). For large diverse markets like the United States, there is also interest in the converse question: why do consumption differences persist (Hart and Alston, 2019). Finally, in addition to measuring historical trends, there has been some research for the United States that has presented forecasts of future alcohol consumption patterns to test whether further convergence in alcohol consumption patterns is likely (Fogarty and Voon, 2018). The extension of the convergence literature to consider future possible consumption patterns raises the question of how alcohol consumption forecasts should be developed. The Autoregressive Integrated Moving Average Model (ARIMA) approach used in Fogarty and Voon is just one possible method for developing alcohol consumption forecasts, and there are strong reasons to suspect that averaging across different

The authors thank Karl Storchmann, an anonymous referee, and the editorial team at JWE for their assistance with this Note.
[a] Agricultural and Resource Economics, School of Agriculture and Environment, University of Western Australia, Crawley WA 6009, Australia; e-mail: derby.voon@uwa.edu.au.
[b] Agricultural and Resource Economics, School of Agriculture and Environment, University of Western Australia, Crawley WA 6009, Australia; e-mail: james.fogarty@uwa.edu.au.

forecasts will lead to an improvement in forecast accuracy (Bates and Granger, 1969; Clemen, 1989). Further, it is also possible to think of alcohol consumption data as hierarchical time series data, where beer, wine, and spirit consumption must add up to total alcohol consumption. For hierarchical time series data, it is possible to impose additional restrictions to ensure consistency in adding up.

In this note we: (i) show how the **R** software platform can be used to obtain alcohol consumption forecasts using a range of different methods and (ii) show that no single forecast approach dominates other methods in terms of forecast performance. To illustrate each method, we use the LaVallee, Kim, and Yi (2014) per capita state level consumption data for the United States.

II. Comparison Setup

The estimation approaches considered are: (i) single equation ARIMA (Box–Jenkins) models (Box et al., 2015); (ii) hierarchical ARIMA models (Hyndman et al., 2011); (iii) single equation state space models (exponential smoothing family) (Hyndman et al., 2008); (iv) hierarchical state space models (Hyndman et al., 2011); (v) the BATS model of De Livera, Hyndman, and Snyder (2011), which extends traditional state space models to allow for complex seasonality through the introduction of a Box–Cox transformation and ARMA errors;[1] and (vi) a neural network model of the form detailed in Hyndman and Athanasopoulos (2018, Ch. 11). For estimation we rely on two **R** packages: Hyndman (2017) and Hyndman et al. (2018).

In this application our focus is to compare the performance of different forecast approaches, and so we separate the data set into a training set (1970 to 2007) and a test set (2008 to 2012). For each type of forecasting method we choose the model form that minimizes AIC, over the training set, and then compare model performance using RMSE across the test set. Figure 1 provides an overview of how to apply each forecasting method in **R**, and the supplementary material provides complete worked examples for each forecast method listed in Figure 1.

III. Results

The first approach used to compare forecast method performance is a series of violin plots of RMSE values, where RMSE values are grouped by estimation method and beverage type. The take home messages from Figure 2 are: (i) in terms of RMSE, relatively simple forecast models perform at least as well as more complex models;

[1] The acronym reflects: (i) the Box–Cox transformation of the time series, (ii) the inclusion of ARMA errors, and (iii) the inclusion of Trend, and Seasonal components.

https://doi.org/10.1017/jwe.2019.15 Published online by Cambridge University Press

Figure 1

Forecasting Alcohol Consumption with *R*

Detail	*Code Details*
	Single equation models
Load data	`my.data <- read.csv('Raw.csv')`
Create time series	`data <- ts(my.data[,1], start=start.year, end=end.year)`
Define training set	`train.data <- window(data, start=start.year, end=end.train.year)`
Define test set	`test.data <- window(data, start=end.train.year+1, end=end.year)`
Set forecast period	`forecast.length <- test.data`
Box–Jenkins	`forecast(auto.arima(train.data, h=forecast.length)`
State space	`forecast(ets(train.data, h=forecast.length)`
Neural network	`forecast(nnetar(train.data, h=forecast.length)`
BATS	`forecast(tbats(train.data, h=forecast.length)`
	Joint equation models
Create time series (a)	`bws <- ts(my.data[,1:3], start=start.year, end=end.year)`
Create time series (b)	`my.dat <- hts(bws, bnames = colnames(bws))`
Define training set	`train.data <- window(my.dat, start=start.year, end=end.year)`
Box–Jenkins	`forecast(train.data, h=forecast.length, fmethod='arima'),level =1)`
State space	`forecast(train.data, h=forecast.length, fmethod='ets'),level =1)`

Figure 2

Violin Plots Comparing Model Performance

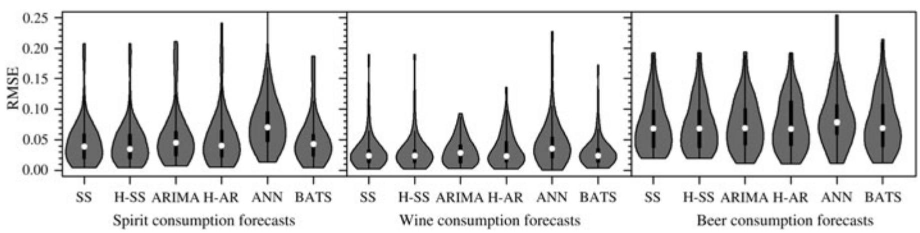

Note: SS = single equation state-space; H-SS = hierarchical single equation state-space; ARIMA = single equation ARIMA;H-AR = hierarchical ARIMA; ANN = autoregressive neural network; BATS = BATS.

(ii) for a given beverage type, forecast models appear to have similar performance; and (iii) across all forecast model types, wine forecasts tend to be the most accurate and beer forecasts least accurate.

Figure 3 plots the maximum and minimum RMSE value for each state by beverage combination across the six forecast methods, and the plots show that there is significant variation in forecast performance between methods across the various state by beverage combinations. Although the plots place the variation in model performance in perspective, they do not show whether one forecast method systematically out performs, when forecasting future alcohol consumption.

To provide a measure of the relative performance of each forecast method, for each state and beverage, the method with the lowest RMSE was identified, and the information is summarized in Table 1. As can be seen from Table 1, at the individual beverage level, the approach that, on most occasions, minimized RMSE,

Figure 3

Within State Variation in Model Performance: RMSE Comparison

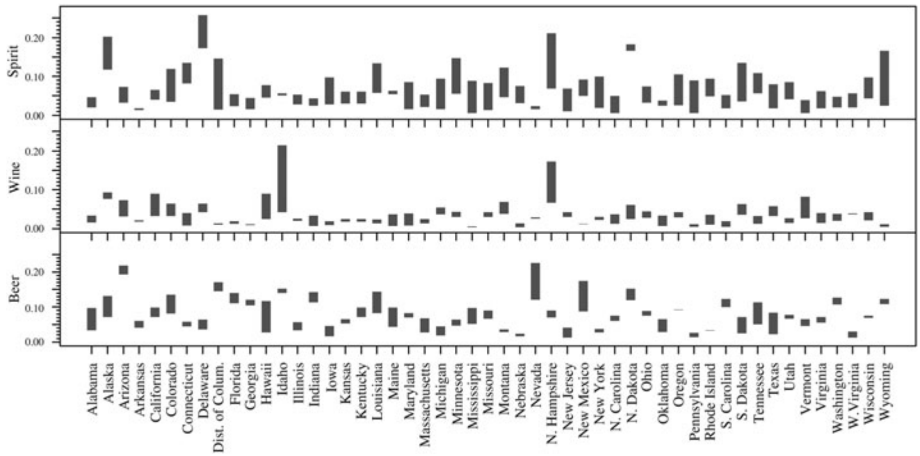

Table 1

Convergence Measures: Coefficient of Variation and Trace

Model	*Spirits (No.)*	*Wine (No.)*	*Beer (No.)*	*Share (%)*
State space	7.5	7.5	11.5	17.3
ARIMA	8	15	6	22.2
Neural network	5	10	13	13.7
BATS	13	7	8	21.6
Hierarchical SS	8.5	7.5	3.5	12.7
Hierarchical ARIMA	9	5	5	12.4

Note: If tied each method allocated 0.5.

varied with beverage type. For spirits the forecast method that most often minimized RSME was BATS; for wine it was the single equation ARIMA method; and for beer it was the autoregressive neural network method. As can be seen from the final column of Table 1, performance across forecast methods, in terms of minimizing RMSE, is quite similar. We do not place special emphasis on particular threshold values for Type I errors, and so with $p = 0.06$ for a proportions test of equality across methods, we simply conclude that there is no strong evidence that one specific forecast method systematically outperforms another, when forecasting future alcohol consumption.

To understand the extent of the differences in forecast consumption levels for each beverage type across methods, the five-year-out forecast values for each method, along with the actual value (black dash), are plotted in Figure 4. As can be seen, there is considerable variation in the forecast level of consumption, across models,

Figure 4

Five Year Future Forecast Comparison: Per Capita Ethanol Consumption

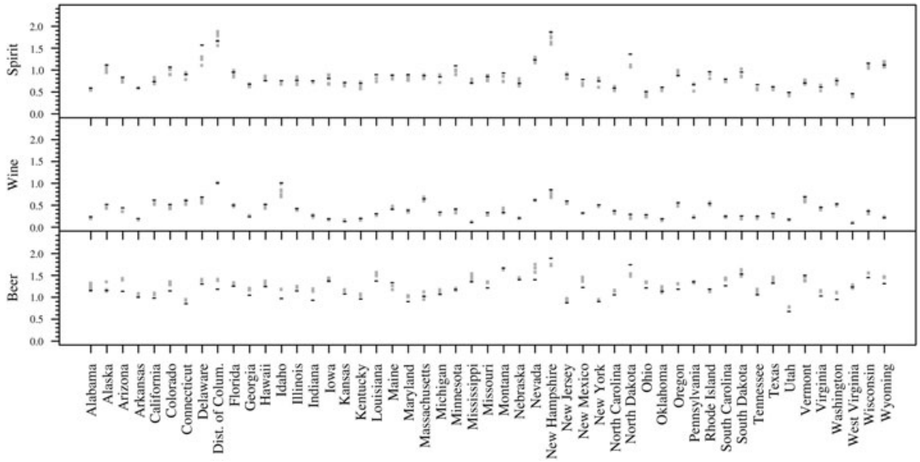

and presenting these different forecasts can help place forecast uncertainty in perspective. For the five-year-out forecast values, the average difference between the maximum forecast value and the minimum forecast value across methods, expressed as a percentage of the actual consumption level, was 12.5% for spirits (SD 6.4%), 11.1% for wine (SD 6.9%), and 5.1% for beer (SD 4.1). So the variation in long range forecasts across methods is non-trivial.

IV. Conclusion

Forecasting future alcohol consumption values allows a range of interesting hypotheses to be considered. Recent work published in the *Journal of Wine Economics* has focused on developing alcohol consumption forecasts using the ARIMA method. In this Note we highlight a range of alternative forecast methods that perform at least as well as the ARIMA method, and show how these methods can be implemented in **R**. To facilitate the use of these methods, a worked example file is provided as part of the supplementary material.

Supplementary Material

For supplementary material accompanying this paper visit https://doi.org/10.1017/jwe.2019.15.

References

Aizenman, J., and Brooks, E. (2008). Globalization and taste convergence: The cases of wine and beer. *Review of International Economics*, 16(2), 217–233.

Bates, J. M., and Granger, C. W. (1969). The combination of forecasts. *Journal of the Operational Research Society*, 20(4), 451–468.

Box, G. E., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control*. Hoboken, NJ: John Wiley & Sons.

Clemen, R. T. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4), 559–583.

Colen, L., and Swinnen, J. (2016). Economic growth, globalisation and beer consumption. *Journal of Agricultural Economics*, 67(1), 186–207.

De Livera, A. M., Hyndman, R. J., and Snyder, R. D. (2011). Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association*, 106(496), 1513–1527.

Fogarty, J., and Voon, D. (2018). Alcohol consumption in the United States: Past, present, and future trends. *Journal of Wine Economics*, 13(2), 121–143.

Hart, J., and Alston, J. M. (2019). Evolving consumption patterns in the U.S. alcohol market: Disaggregated spatial analysis. *Journal of Wine Economics*, forthcoming.

Holmes, A. J., and Anderson, K. (2017). Convergence in national alcohol consumption patterns: New global indicators. *Journal of Wine Economics*, 12(2), 117–148.

Hyndman, R., Koehler, A. B., Ord, J. K., and Snyder, R. D. (2008). *Forecasting with Exponential Smoothing: The State Space Approach*. Berlin, Germany: Springer Science & Business Media.

Hyndman, R., Lee, A., Wang, E., and Wickramasuriya, S. (2018). hts: Hierarchical and grouped time series. R package version 5.1.5.

Hyndman, R. J. (2017). forecast: Forecasting functions for time series and linear models. R package version 8.2.

Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., and Shang, H. L. (2011). Optimal combination forecasts for hierarchical time series. *Computational Statistics & Data Analysis*, 55(9), 2579–2589.

Hyndman, R. J., and Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. OTexts.

LaVallee, R. A., Kim, T., and Yi, H.-y. (2014). Surveillance report# 98: Apparent per capita alcohol consumption: National, state, and regional trends, 1977–2012. US Department of Health and Human Services (DHHS).

Mills, T. C. (2018). Is there convergence in national alcohol consumption patterns? Evidence from a compositional time series approach. *Journal of Wine Economics*, 13(1), 92–98.

Smith, D. E., and Mitry, D. J. (2007). Cultural convergence: Consumer behavioral changes in the European wine market. *Journal of Wine Research*, 18(2), 107–112.