



HAL
open science

Forecasting Bordeaux wine prices using state-space methods

Stephen Bazen, Jean-Marie Cardebat

► **To cite this version:**

Stephen Bazen, Jean-Marie Cardebat. Forecasting Bordeaux wine prices using state-space methods. Applied Economics, 2018, 50 (47), pp.5110 - 5121. 10.1080/00036846.2018.1472740 . hal-01867216

HAL Id: hal-01867216

<https://amu.hal.science/hal-01867216>

Submitted on 8 Apr 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Forecasting Bordeaux wine prices using state-space methods

Stephen Bazen^a and Jean-Marie Cardebat^{b,c}

^aAix-Marseille University (AMSE), CNRS and EHESS, Marseille, France; ^bLAREFI, Université de Bordeaux, Bordeaux, France; ^cINSEEC, Bordeaux, France

ABSTRACT

Generic Bordeaux red wine (basic claret) can be regarded as being similar to an agricultural commodity. Production volumes are substantial, they are traded at high frequency and the quality of the product is relatively homogeneous. Unlike other commodities and the top-end wines (which represent only 3% of the traded volume), there is no futures market for generic Bordeaux wine. Reliable forecasts of prices can to large extent replace this information deficiency and improve the functioning of the market. We use state-space methods with monthly data to obtain a univariate forecasting model for the average price. The estimates highlight the stochastic trend and the seasonality present in the evolution of the price over the period 1999 to 2016. The model predicts the path of wine prices out of sample reasonably well, suggesting that this approach is useful for making reasonably accurate forecasts of future price movements.

KEYWORDS

Wine prices; state-space methods; forecasting

JEL CLASSIFICATION

C53; L66; Q11

I. Introduction

Price forecasting in commodity markets plays a crucial role in market participants' decision-making. Investment plans and buying and selling decisions in these markets are based on price expectations. A huge literature exists on forecasting commodity prices. The different approaches adopted include identifying the underlying fundamentals that determine commodity prices such as storage and demand variables (Williams and Wright 2005); forecasting price movements on the basis of futures prices (Fama and French 1987; Tomek 1997); basing commodity price forecasts on exchange rates (Chen, Rogoff, and Rossi 2010); using co-movements across commodity prices to forecast individual commodity prices (West and Wong 2014); applying neural networks (Kohzadi et al. 1996); and estimating some form of time series model such as GARCH (Wei, Wang, and Huang 2010).

Wine however appears to be an exception: there is a signal lack of prediction tools for a sector which is worth more than 300 billion USD in 2015 globally.¹ Wine-makers, sellers, buyers and investors, all seek reliable price forecasts. However, the wine economics literature

does not really address this issue. Ashenfelter (2008) and Oczkowski (2010) have attempted to forecast fine wine prices based on hedonic regressions. This approach allows predictions of the price of given wine at time t according to its characteristics. However, the hedonic regression by its very nature is not the most suitable method for price forecasting (i.e. a price prediction for $t + n$). More recently, Yeo, Fletcher, and Shawe-Taylor (2015) have used machine-learning methods for predicting *returns* on fine wines, thus treating them as investments. However, these few approaches have been confined to wines at the top-end of the market.

However, these top-end wines represent no more than 3% of the total production of Bordeaux wine. Generic Bordeaux red wine (basic claret) represents about 40% or 2 million hectolitres of the total production. Bordeaux generic wine can be considered to large extent in the same way as an agricultural commodity in the sense that large volumes are traded at a high frequency, and the quality is fairly constant over time. However, unlike most agricultural commodities, there is no futures market.² There is no 'future price' for wine: market participants do not possess the information necessary to form expectations.

CONTACT Jean-Marie Cardebat ✉ jean-marie.cardebat@u-bordeaux.fr 📍 Avenue Léon Duguit 33600, Pessac, France

¹Mordor Intelligence, Global Wine Market – Growth, Trends and Forecasts (2017–2022), 2017.

²For top-end wines, there is the *primeur* market which involves purchasing the latest vintage, while is still aging in the barrel and before it is bottled. Based on tasting from the barrel, a price is struck for delivery at some later date.

Forecasting future wine prices can therefore play a role in providing agents in the Bordeaux generic wine market with information permitting them to form expectations of price movements which improve the management of stocks, and thereby increase market efficiency. Prices are subject to quite large variations and create uncertainty for the actors in this market with regard to buying and selling decisions, and so there is a useful role for price forecasting. The purpose of this paper is to fill this gap by using a state-space approach in order to provide forecasts of the price of generic Bordeaux red wine.

One reason for why this has not already been undertaken is the lack of availability of appropriate data. We obtained monthly data for price, volume and the number of contracts, along with annual harvest quantities directly from the Bordeaux wine trade professional organization (CIVB), for the period 1999 to 2016. Using state-space methods (see e.g. Harvey 1989) to obtain forecasts of Bordeaux generic wine prices, we estimate a univariate model for the average price of wine. The estimated model is used to make in sample and out of sample predictions. The quality of the predictions is generally good. In cases when the predictions may be poor for a given period, future predictions are brought back on track due to the efficiency of updating mechanism of the Kalman filter.

The interest of this paper is then threefold. First, it fills the gap in the wine economics literature where the price of generic wines has been largely ignored; however, these wines represent largest the traded volume. The issue of price forecasting has received virtually no attention unlike other commodities. Second, the paper contributes to the price forecasting literature in the commodity sector. Third, the quality of the predictions made with this approach could be useful for the wine professionals and help them in forming their expectations. The paper begins by setting out the market context and reviewing the literature on the determinants of wine prices in [Section II](#). The data and the econometric approach used are described in [Section III](#), and the estimates and forecasting properties are examined in [Section IV](#). Brief conclusions are presented in [Section V](#).

II. Market context and existing literature

According to the Observatoire International du Vin (OIV), the Bordeaux wine region is of interest because it is the biggest wine-producing area in the world.³ Since the 2000s, fine wine has been increasingly treated as a financial asset (Masset and Henderson 2010) and a large strand of the wine economics literature has focused on the determinants of Bordeaux wine prices at the top end of the market. The aim is usually to identify the major influences using the hedonic price approach (see e.g. Cardebat and Figuet, 2004). However, these top-end wines represent no more than 3% of the total production of Bordeaux wine. The Bordeaux region comprises 57 ‘appellations’ or AOCs and the prominent wine type among these is the Bordeaux AOC or generic (basic claret). This AOC represents about 40% or 2 million hectolitres of the total production of Bordeaux wines.

Like the top-end wines, the price of Bordeaux generic AOC wine has also increased since the mid 2000s, following a period of decline and stagnation from 1997 to 2004 (see [Figure 1](#)). The wholesale claret market was dogged by overproduction throughout the 1990s and up to the mid 2000s. The ‘Bordeaux plan’ was agreed between producers and the Ministry of Agriculture, and implemented in 2006. The aim was to bring an end to overproduction through a massive reduction in growing capacity (in which vineyards were destroyed). At the same time demand has expanded: while French and continental European demand has been declining for more than 30 years, there was a substantial increase in demand from Asia. The combination of reduced production and higher demand led to a significant turnaround in prices. Weak harvests at the beginning of the 2010s in Bordeaux also contributed to a tighter market and further fuelled increases in price.

This increase in price occurred in spite of more intensive competition from foreign producers. Spain is the world leader in both wholesale wine production and in export terms, as well as prices are three times lower than those in France and Italy (based on OIV data, reference in footnote 3). While the price of Bordeaux red wine has remained buoyant, its market shares have been declining. Australian wines increased their market shares to the detriment of Bordeaux red wines in the US, UK

³See OIV, <http://www.oiv.int/fr/bases-de-donnees-et-statistiques>.

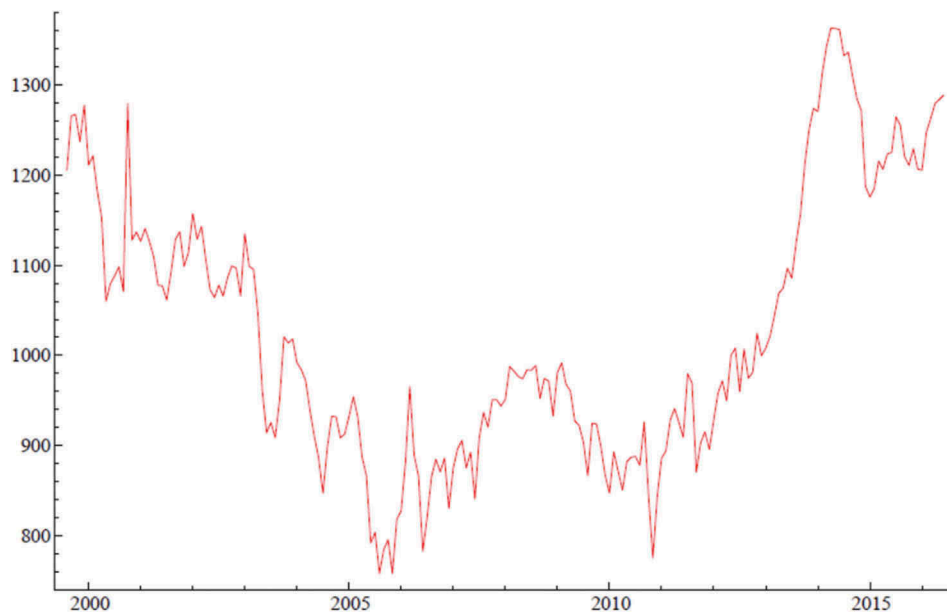


Figure 1. Monthly average price of generic Bordeaux red wine.

and Chinese markets, and have seen the fastest growth in exports. Since 2000 Chile has emerged as a major player and is now in fourth place behind France, Italy and Spain.

The functioning of the Bordeaux generic wine market is characterized by the risk aversion of producers who hold limited stocks (the average is usually comprised between 0.5 and 1 times the harvest, based on CIVB data) and decide on how much wine to put on the market while seeking to avoid being unable to supply the market should there be a bad harvest in the following year. For instance, a buyer (who must first be registered on the market at CIVB) wishes to purchase a quantity of claret to be sold under its own label and approaches an intermediary: the offer is opened up to producers who will sell some of their existing stock depending on the price and their need for cash revenue. However, it will also depend on the extent of remaining stock, since the producer will always desire to hold some wine back in case the next harvest is not a good one in volume terms.⁴

The next harvest will depend on climatic factors which affect the survival of the grape on the vine once the latter flowers. Through the spring frost may kill the grape flower and during the following months hail storms may damage the grapes.

General precipitation and lack thereof at particular points in time through the cultivation period will also be relevant factors, and there could be problems due to certain diseases (Jones and Davis 2000). Taken together, the decision to sell part of existing stocks will be more likely to occur once the producer has less uncertainty concerning the size of forthcoming harvest. Thus, there will be more sales in volume terms after the period of risk of frost is over, and then again when the nature of the summer weather has been ascertained. There is also the matter of freeing up space for stocking the new harvest. This suggests that over a 12-month period there will be variable supply giving rise to seasonality in prices.

The wine economics literature focuses on price determinants and price dynamics, but not on price forecasting. A recent paper used machine-learning methods to predict returns on fine wines (Yeo, Fletcher, and Shawe-Taylor 2015). Here, we are concerned by the price rather than the return since generic wine cannot be laid down and re-sold. The storage period cannot exceed 3 years for a generic wine without a deterioration in quality. The determinants of wine prices, usually modelled by hedonic regressions, are the quality signals revealed on the label or/and based on the opinions of wine experts. The quality signals on the labels are manifold. The

⁴While there are three recognized quality levels within the AOC, since the label only refers to the 'terroir' and composition of grape varieties, quality plays only a small role in this particular market.

vintage is an indicator of quality for people with substantial knowledge (the vintages are more or less good depending on the weather, especially in European wine producing-countries). Official rankings (like Grands Crus classés in Bordeaux region) are a well-known signal of quality rooted in a long history (from 1855 for the Médoc ranking for instance). Obviously, the brand and the denomination of origin (DO) are also major determinants of the price. Both are reputation variables. The brand name is a signal of individual reputation, whereas the DO is a signal of collective reputation, but both affect the price (Costanigro, McCluskey, and Goemans 2010; Schamel 2009; Frick and Simmons 2013).

Experts' scores and wine trade competition medals or awards are also important quality signals, especially Robert Parker's grades (see Ali, Lecocq, and Visser 2008). There is a consensus in the literature on their significant positive impact on wine prices (for a survey see Storchmann 2012), although these scores are often biased (Cardebat, Figuet, and Paroissien 2014; Oczkowski 2016) and sometimes dubious (Hodgson 2008). One interesting aspect of this literature about wine experts is that climatic conditions during the growing and the harvest seasons are also a major predictor of the scores (Ashenfelter 2008; Cardebat and Livat 2016) and of the price of wine (Jones and Storchmann 2001; Lecocq and Visser 2006; Dubois and Nauges 2010; Cardebat, Figuet, and Paroissien 2014).

The majority of these determinants make sense only for fine wines obviously, especially experts' grades, or ranking. However meteorological conditions also have an impact on the price of generic wines where the impact is measured in terms of volume (whereas for fine wines it is viewed more in terms of quality) and thereby prices in the generic wine market. In fact, this segment of the wine market can be considered in the same way as agricultural production where volume appears as the driver of the price. Even though the wine economics literature has not yet explored this aspect of the wine market, it would be interesting to note that generic wine prices could be driven by strategic storage behaviour. Such an explanation is traditional in agricultural literature since the seminal paper of Gustafson (1958) and well illustrated by Williams and Wright (2005). Gouel (2012) also offers a synthetic presentation of the different storage models developed in the

agricultural field, especially when exogenous shocks (like unfavourable climatic conditions) affect the supply. While these factors will be relevant for the overall trend in wine prices, they may be less relevant for explaining the short-run movements found in monthly price series.

III. Data and model specification

The specific series that we use is the monthly average price per barrel (containing 900 l) for the period late 1999 to June 2016. The average is taken of the prices in contracts signed in the period. Information is also available on the volume sold and the number of contracts signed. The price series (Figure 1) shows a trend decline between 2000 and 2005, then rises slightly before falling back to its 2005 level in 2010. Thereafter the price rises sharply, almost exponentially, to a record in early 2014 before falling back slightly over the remainder of 2014 and remaining in a limited range. It is interesting to note that the volumes of sales do not vary in the same way (see Figure 2) and on their own cannot explain why the price increases dramatically after 2010. The role of the grape harvest will be factor since there are at least two disappointing harvests in the period examined (in 2008 and 2013; see Figure 3).

The price series is analysed using a state-space model which is estimated for the purposes of predicting wine prices. This kind of model is based on the notion that a series can be decomposed into trend, cycle, seasonal and irregular components. The first three constitute a pattern of regularity in the data, but each of the components is inherently unobservable. Sometimes certain variables can be used as proxies such as a GDP measure for the cyclical component, but this often proves unsatisfactory. The different components can be replaced by artificial scales (linear trends and dummy variables), resulting in a deterministic representation of these components. In the state-space approach, pioneered in econometric modelling by Harvey (1989), and set out very clearly in an accessible form for practitioners in Commandeur and Koopman (2007), these components are treated as unobservable and stochastic. The deterministic formulation is in fact a special case in these models. From a forecasting point of view, the flexibility of the state-space approach allows previous forecast errors to be




Figure 2. Annual harvest of generic Bordeaux red wine.




Figure 3. Monthly sales volumes.

rapidly taken into account in the modelling of the different components so that future forecasts will be more accurate. Such an approach is both a useful benchmark for evaluating richer models and also has the advantage of requiring no more data than on the series to be predicted.

The series for wine prices over the period covered potentially contains stochastic trend, seasonal and irregular components⁵:

$$y_t = \mu_t + \gamma_t + \varepsilon_t \quad \varepsilon_t \sim \text{NID}(0, \sigma_\varepsilon^2) \quad (1)$$

where ε_t is the irregular term and μ_t is the stochastic trend determined as

$$\mu_t = \mu_{t-1} + \beta_{t-1} + \eta_t \quad (2)$$

where β_t is the slope of the trend and is also stochastic:

$$\beta_t = \beta_{t-1} + \zeta_t \quad (3)$$

The stochastic terms η_t and ζ_t are assumed to have zero means and their variances are σ_η^2 and σ_ζ^2 , respectively.

In fact, it is assumed that $\eta_t \sim \text{NID}(0, \sigma_\eta^2)$ and $\zeta_t \sim \text{NID}(0, \sigma_\zeta^2)$. If the estimated variance of either of these components is close to zero, then that component can be treated as being deterministic. If $\sigma_\zeta^2 = 0$

⁵A cyclical component is not included since the time period is short, and any cyclical movements will be subsumed into the trend and seasonal components.

then the slope β is constant across the sample period since $\beta_t = \beta_{t-1}$. If $\sigma_\eta^2 = 0$ then the trend is deterministic. The seasonal component, γ_t , can be modelled using dummies plus a stochastic term for each month $j = 1, 2, \dots, 12$:

$$\gamma_{jt} = \gamma_{jt,-1} + \omega_{jt} \quad \omega_{jt} \sim \text{NID}(0, \sigma_\omega^2)$$

Only one of these components affects the dependent variable at a given time and so:

$$\gamma_t = \gamma_{jt} \text{ for } t = 1, 2, \dots, T$$

If ω_{jt} has a zero variance ($\sigma_\omega^2 = 0$) then the seasonality is deterministic.

The different parameters of the model are obtained using Gaussian state-space methods, involving the Kalman filter and maximum-likelihood estimation. The estimates of the variances of the different stochastic terms and the predictions presented here were obtained using version 8 of the STAMP software package (Koopman et al. 2007).

IV. Model estimates and forecasting properties

The first stage of the approach is to estimate the hyperparameters, i.e. the variances of the different components. If the estimate of any of these is close to zero, then the component can be treated as deterministic. The estimated variances of the slope and seasonal components are small relative to the variance of the irregular term (see Table 1). In this basic specification, there is some evidence of mild autocorrelation among the residuals and the normality test fails. Given the dynamics embedded in the specification of the components – for example, the trend is specified as a random walk with drift – and the way in which seasonality is modelled, the inclusion of lagged values in this kind of model is not always necessary. However the addition of 12 lags removes the autocorrelation, although only the first, ninth, tenth and twelfth lags are statistically significant at conventional levels. There is no indication of the presence of heteroscedasticity, and the normality test is satisfactory.

The second stage arises out of the estimation process, in which an estimate of the different components is made at each point in time through the sample period and these are presented in Figure 4. The trend or level component is clearly stochastic

and tracks the actual price of wine quite closely. It would be difficult to model the path of wine prices with a deterministic trend specification. The stochastic nature of seasonal component is apparent from the fact that the seasonal variation in prices is less pronounced after 2006. This feature of price variation links in with the nature of the behaviour of Bordeaux wine producers mentioned in the introduction, where uncertainty about the size of the grape harvest is attenuated or exacerbated at certain points in the year. However, the impact on price is not that substantial relative to the trend component. The irregular term picks up the remainder of the variation in prices.

This univariate model has no explanatory content although the movements in the trend and seasonal components can be linked to underlying economic factors such as overseas demand and size of harvest. The major advantage of this approach is the possibility of producing forecasts without having to predict the values of any explanatory variables or to define a scenario for the future. However, in any exercise in forecasting there are unpredictable factors which will undermine the accuracy of the forecasts. Furthermore, pure forecasting, i.e. of months and years to come, is unverifiable at the time of making the predictions. Thus in order to assess the adequacy of the model, we examine the extent to which it predicts the trajectory of wine prices in two ways. The first is to examine the *one-step ahead predictions* of the model within the sample period and compare the predicted and observed values. The second is to make *unconditional forecasts* by truncating the sample used to estimate the model parameters and state vector at certain points in time, and then forecast say 6 months ahead out of sample. The latter can be compared to the observed outcomes as if one were forecasting the future path of wine prices. It is unlikely that forecasts over a longer horizon than 6 months are likely to be meaningful in view of the stochastic nature of the trend.


The one-step predictions are made for the period from January 2014 to the end of the sample period (see Figure 5). This within sample predictive performance is satisfactory in statistical terms, in the sense that the predictions and the observed values both inside the 90% confidence interval except for 1 month (October 2014) and only at this point does the CUSUM plot lie outside the confidence interval. It

Table 1. Estimated component variances and diagnostic tests.

Dependent variable: price of Bordeaux AOC red wine	No lagged dependent variables	Twelve lags on the dependent variable
Estimated component variances		
Trend	1309.35	99.54
Slope	0.330	0.064
Seasonal	0.609	1.433
Irregular	206.57	577.28
Diagnostic tests* (5% critical value in parentheses)		
Autocorrelation (24th order)	38.81 ($\chi^2_{21} = 32.67$)	15.239 ($\chi^2_{21} = 32.67$)
Heteroscedasticity	0.86 ($F(64, 64) = 1.51$)	0.680 ($F(56, 56) = 1.56$)
Normality	13.25 ($\chi^2_2 = 5.99$)	1.839 ($\chi^2_2 = 5.99$)
Number of observations	205	193

*See Harvey (1989) for descriptions of these tests (pp. 258–260).

Trend or level (in red, actual price in black)



Seasonal and irregular components





Figure 4. Estimated model components.

One step predictions



CUSUM plot




Figure 5. One-step predicted and actual prices and CUSUM plot for 2014:1 to 2016:VI (with 90% confidence intervals) – blue line is the prediction and red line is the observed value.

is important to note that whenever there is a prediction error, the one-step predictions quickly come back onto track and CUSUM plot begins to return towards zero. This capacity to rapidly correct prediction errors underlines the efficiency of state-space methods for forecasting purposes. Unconditional price forecasts are made for the period January to June 2016 and are grouped around the observed prices (see Figure 6). The overall forecasting ability of the model for the latter end of the sample period

would appear to be highly satisfactory and is superior to the predictions from a simple random walk (given by the dotted line in the figure).

In order to assess the forecasting performance of the approach in a more general fashion, we examine forecast accuracy for a certain number of sub-periods. The monthly average price of Bordeaux generic wine reaches its maximum value for the sample period in April 2014. In Figure 7, four consecutive 6-month sub-periods




Figure 6. Unconditional forecasts 2016:I to 2016:VI (with one standard error confidence intervals) – red line is the forecast and blue dots are the observed price.




Figure 7. Sub-periods for forecast evaluation.

are defined beginning with January 2014. The model is estimated up to the month prior to the beginning of each sub-period, and unconditional forecasts are made for the following 6-month period. In this way, the ability of the model to provide accurate forecasts during the period when the wine price peaked and then fell away quite dramatically before stabilizing can be assessed. The outcomes are presented in [Figure 8](#).

In three of the four sub-periods, the forecasting performance is highly satisfactory and generally superior to predictions from a simple random walk. However, the model fails to predict the extent of the decrease in wine prices in the second half of 2014: the prediction for July 2014 is already on the lower limit of the confidence interval, and the forecast error gets progressively worse thereafter. However, when the experience of 2014 is integrated




Figure 8. Unconditional forecasts for various sub-periods (with one standard error confidence intervals) – red line is the forecast and blue dots are the observed prices.

in the model, the forecasts for 2015 are far more satisfactory and improve as the sample period is extended through 2015. This exercise shows both the strengths and weaknesses of a univariate state-space model. Additional sample information is used efficiently to rapidly correct prediction errors. This is due to the way in which the model is estimated using the error decomposition form of the log likelihood and the Kalman filter. But there is little in the model that is capable of predicting a major turning point in a series and the subsequent depth of negative price shock. A multivariate or structural model would in principle be better in this respect. However, there is a trade-off in terms of the information requirements when producing forecasts: no information about the future is required in a univariate model. In a multivariate or structural model, the future paths of explanatory variables and/or the elaboration of a relevant scenario for the future in terms of exchange rate policy, environmental factors or the state of the

economies where generic Bordeaux red wine is purchased must be defined.

V. Conclusions

Due to their flexibility and ability to adapt to new information, state-space methods have a number of strengths in forecasting the path of time series variables. In the context of generic Bordeaux red wine prices, these methods highlight the stochastic trend and the seasonality present in the evolution of the price over the period 2000 to 2016. They also satisfactorily predict the path of wine prices out of sample, suggesting that the models could be useful for making accurate forecasts of future price trends. In the absence of futures markets for wines at the lower end of the market, such forecasts could in principle provide additional information to market participants and improve the functioning of the bulk market for generic wines. This paper therefore contributes to both the wine

economics and commodity price forecasting literatures, as well as providing a tool for wine professionals.

This work could be extended in many directions. First, it would be obviously interesting to make forecasts for other large bulk wine-producing regions in Spain, Chile or Australia for instance. Second, other methods could be used to make forecasts of the price of Bordeaux red wine. A structural approach based on competitive storage models could be promising and a comparison of the results obtained from the two approaches could be of interest. The assumption would be that state-space methods are useful for short term forecasts while storage models, which integrate weather conditions and harvest predictions, would be relevant for medium term forecasts. The methods could be complementary. Third, this paper raises the issue for professionals of creating a future market for generic wine, which would make sense in a globalized sector. What could be the conditions for establishing such a market? What would be traded: varietal wines such as Merlot or Cabernet, or geographical delimited wines like Bordeaux or Rioja?

Acknowledgments

An earlier version of this paper was presented at the American Association of Wine Economics Conference, Bordeaux, 21–25 June 2016.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Ali, H. H., S. Lecocq, and M. Visser. 2008. "The Impact of Gurus: Parker Grades and En Primeur Wine Prices." *The Economic Journal* 118 (529): F158–F173.
- Ashenfelter, O. 2008. "Predicting the Quality and Prices of Bordeaux Wine." *The Economic Journal* 118 (529): F174–F184.
- Cardebat, J. M., and J. M. Figuet. 2004. "What Explains Bordeaux Wine Prices?" *Applied Economics Letters* 11 (5): 293–296.
- Cardebat, J. M., J. M. Figuet, and E. Paroissien. 2014. "Expert Opinion and Bordeaux Wine Prices: An Attempt to Correct Biases in Subjective Judgments." *Journal of Wine Economics* 9 (3): 282–303.
- Cardebat, J. M., and F. Livat. 2016. "Wine Experts' Rating: A Matter of Taste?" *International Journal of Wine Business Research* 28 (1): 43–58.
- Chen, Y. C., K. S. Rogoff, and B. Rossi. 2010. "Can Exchange Rates Forecast Commodity Prices?" *Quarterly Journal of Economics* 125 (3): 1145–1194.
- Commandeur, J., and S. J. Koopman. 2007. *An Introduction to State Space Time Series Analysis*. Oxford: Oxford University Press.
- Costanigro, M., J. J. McCluskey, and C. Goemans. 2010. "The Economics of Nested Names: Name Specificity, Reputations, and Price Premia." *American Journal of Agricultural Economics* 92 (5): 1339–1350.
- Dubois, P., and C. Nauges. 2010. "Identifying the Effect of Unobserved Quality and Expert Reviews in the Pricing of Experience Goods: Empirical Application on Bordeaux Wine." *International Journal of Industrial Organization* 28 (3): 205–212.
- Fama, E. F., and K. R. French. 1987. "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage." *Journal of Business* 60: 55–73.
- Frick, B., and R. Simmons. 2013. "The Impact of Individual and Collective Reputation on Wine Prices: Empirical Evidence from the Mosel Valley." *Journal of Business Economics* 83 (2): 101–119.
- Gouel, C. 2012. "Agricultural Price Instability: A Survey of Competing Explanations and Remedies." *Journal of Economic Surveys* 26 (1): 129–156.
- Gustafson, R. L. 1958. "Implications of Recent Research on Optimal Storage Rules." *Journal of Farm Economics* 40 (2): 290–300.
- Harvey, A. 1989. *Forecasting, Structural Times Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Hodgson, R. T. 2008. "An Examination of Judge Reliability at a Major US Wine Competition." *Journal of Wine Economics* 3 (2): 105–113.
- Jones, G. V., and R. E. Davis. 2000. "Climate Influences on Grapevine Phenology, Grape Composition, and Wine Production and Quality for Bordeaux, France." *American Journal of Enology and Viticulture* 51 (3): 249–261.
- Jones, G. V., and K. H. Storchmann. 2001. "Wine Market Prices and Investment under Uncertainty: An Econometric Model for Bordeaux Crus Classés." *Agricultural Economics* 26 (2): 115–133.
- Kohzadi, N., M. S. Boyd, B. Kermanshahi, and I. Kaastra. 1996. "A Comparison of Artificial Neural Network and Time Series Models for Forecasting Commodity Prices." *Neurocomputing* 10 (2): 169–181.
- Koopman, S. J., A. Harvey, J. Doornik, and N. Shephard. 2007. *STAMP Version 8.0*. London: Timberlake Consultants.
- Lecocq, S., and M. Visser. 2006. "Spatial Variations in Weather Conditions and Wine Prices in Bordeaux." *Journal of Wine Economics* 1 (2): 114–124.
- Masset, P., and C. Henderson. 2010. "Wine as an Alternative Asset Class." *Journal of Wine Economics* 5 (1): 87–118.

- Oczkowski, E. 2010. "Hedonic Wine Price Predictions and Nonnormal Errors." *Agribusiness* 26 (4): 519–535.
- Oczkowski, E. 2016. "Identifying the Effects of Objective and Subjective Quality on Wine Prices." *Journal of Wine Economics* 11 (2): 249–260.
- Schamel, G. 2009. "Dynamic Analysis of Brand and Regional Reputation: The Case of Wine." *Journal of Wine Economics* 4 (1): 62–80.
- Storchmann, K. 2012. "Wine Economics." *Journal of Wine Economics* 7 (1): 1–33.
- Tomek, W. G. 1997. "Commodity Futures Prices as Forecasts." *Review of Agricultural Economics* 19: 23–44.
- Wei, Y., Y. Wang, and D. Huang. 2010. "Forecasting Crude Oil Market Volatility: Further Evidence Using GARCH-class Models." *Energy Economics* 32 (6): 1477–1484.
- West, K. D., and K. F. Wong. 2014. "A Factor Model for Co-Movements of Commodity Prices." *Journal of International Money and Finance* 42: 289–309.
- Williams, J. C., and B. D. Wright. 2005. *Storage and Commodity Markets*. Cambridge: Cambridge University Press.
- Yeo, M., T. Fletcher, and J. Shawe-Taylor. 2015. "Machine Learning in Fine Wine Price Prediction." *Journal of Wine Economics* 10 (2): 151–172.