


ARTICLE

# How sample bias affects the assessment of wine investment returns

Joseph L. Breeden 

Auctionforecast.com, Santa Fe, NM, USA

Email: [breeden@prescientmodels.com](mailto:breeden@prescientmodels.com) or [breeden@auctionforecast.com](mailto:breeden@auctionforecast.com)

## Abstract

Wine investment returns can come from overall market trends or price increases with age. Because of the short wine price histories available, market and maturation effects are difficult to separate. Consequently, researchers often obtain dramatically different estimates of investment returns. We find that data sample bias may be the hidden cause of the disparate estimates. In wine auction data, the sample bias refers to a shift in the distribution of which wines are traded as a function of their age. Such sample bias in panel data sampled across many different wine labels can distort the estimation of price increases versus age and consequently impact the estimation of market trends. This analysis shows that segmenting the analysis such that the data panels contain wine labels with similar trading characteristics can lead to a more stable estimation.

The analysis here looks at data from Bordeaux, Italy, Australia, and California. An Age-Period-Cohort (APC) analysis is applied to data panels from each region. Then the data in each region is segmented by a measure of popularity in order to reduce sampling bias. Data thus segmented is then re-analyzed to demonstrate the difference in estimating price appreciation lifecycles and market trends.

**Keywords:** age-period-cohort models; auction prices; sample bias; survival analysis; wine investment

**JEL Classifications:** C23; D44; G11; G12; Q11

## 1. Introduction

The potential price appreciation in wine investment is contested both among researchers and investors. Of course, a seasoned wine trader might have excellent returns with a simple buy-low-sell-high approach by hunting for outliers at auction. In analyzing investment potential, this research seeks to enhance the measurement of price appreciation with the age of the wine and create a measurement of wine markets that have been normalized for wine aging. A wine label refers to any bottle of wine as identified by brand and vintage. The analysis of a single label, such as the work by Jovanovic (2008) studying 1870 Lafite, cannot distinguish the market and aging

effects on the price at auction. Rather, a broad dataset covering a range of labels observed over many calendar dates is necessary.

This research was motivated by the observation that the distribution of wine labels traded in the first decade since their release is different from the wine labels traded several decades after their release. If all wines from a region are analyzed together, then one would observe that lower-priced wine labels drop out of the data set with age sooner than higher-priced wines. This bias in the data distribution is referred to as *right-censoring* in the survival modeling literature, and when the censoring is selective based upon an unmodeled variable, it can lead to a bias in estimating the function versus age, which is referred to as *frailty* (Balan and Putter, 2020) when the bias is due to an unobserved explanatory variable. In this research, we hypothesize that trading frequency (popularity) can be used as a segmentation variable to eliminate frailty in the estimates.

In a continuation of our previous research into applying Age-Period-Cohort (APC) analysis to wine prices (Breeden and Liang, 2017; Breeden, 2010), we explored the impact of market segmentation on the estimation of wine investment returns. Specifically, this was done by performing separate APC analyses on different wine segmentations, specifically estimating the wine market and price growth with wine age. Analyzing data from Bordeaux, Italy, Australia, and California, we examined how the same segmentation approach performed in each region.

This analysis will show that measurements of the wine market are remarkably stable relative to the segmentation scheme chosen, but price appreciation with age varies strongly depending upon how one defines the segments. With the cause appearing to be due to data censoring for less popular wines, we developed a forward-looking segmentation algorithm that predicts future popularity so that even young wines may be properly segmented in order to reduce bias due to censoring. Further, comparisons across regions not only provide a useful cross-validation of the approach but also provide market and price appreciation with age that are not well studied outside Bordeaux.

This paper proceeds as follows: Section II provides a literature review on wine as an investment. Section III describes the data used in the study. Section IV describes the APC approach for estimating market trends and price growth with the age of a wine. Section V provides all of the results by region studied, and Section VI describes our conclusions.

## II. Wine investment literature

The ultimate goal of this research is to predict future returns on investing in a specific wine or portfolio of wines. The earliest analysis of wine prices was by Ashenfelter (2010). Wood and Anderson (2006) employed a cubic age function and demonstrated that the age of a wine was predictive of price. Various authors have studied the rate of return for holding wines and the diversification benefit of having wine as part of an overall portfolio. See Storchmann (2012) for a detailed survey of the literature.

With a wine-only portfolio, one hopes that wines have a higher rate of return than other investments (Burton and Jacobsen, 2001; Fogarty, 2006; Jaeger, 1981; Krasker, 1979; Jones and Storchmann, 2001), although not all studies have found long-term

positive returns. More recently, Sanning, Shaffer, and Sharratt (2008) show that wine has exhibited large excess returns. However, any study conducted immediately prior to the 2011 wine market correction is likely to be overly optimistic. As with most markets, one must consider long periods in order to capture both rallies and corrections.

With a diversification approach (Bouri et al., 2018; Chu 2014), a lower-return series that is somewhat uncorrelated to other investment vehicles is also beneficial. So, any return modestly greater than zero could be beneficial if uncorrelated or, better yet, anti-correlated to other instruments. Initial studies did find low or negative correlation to typical investment vehicles (Fogarty, 2010; Masset and Henderson, 2010; Masset and Weisskopf, 2018; Sanning, Shaffer, and Sharratt, 2008), but intrinsic return remains a question.

Fogarty and Sadler (2014, table 1) summarize previous studies, appropriately concluding that researchers have obtained a wide variety of results depending upon the time period examined and the method of analysis applied. The preponderance of the literature appears to take a portfolio perspective where one presumably has a diverse collection of wines of many ages so that only the overall market matters. However, if all wine prices increase with age or at least do not decrease throughout the tradable life of a wine, then there is an inherent price growth rate of the portfolio separate from the market. The only way to purely measure market trends for a portfolio of wines would be to steadily replace older labels with newer labels such that the average age remains stable.

### III. Data

The current study leverages wine prices from AuctionForecast.com. Table 1 provides a summary of the date and vintage ranges for the analysis by segment, as well as the number of auction lots and wine-vintages. In all usages, vintage refers to the year the grapes were grown.

The data was collected from Acker Wines, Bidforwine, Bonhams, Christie's, Chicago Wine Company, Hart Davis Hart, Langton's, Skinner, Sotheby's, Spectrum, Veiling Sylvie's, and Zachy's. Only prices from homogenous lots were retained, meaning the lot price pertains to a single wine-vintage and bottle size. The lot price was divided by the number of bottles in the lot, and all prices were

**Table 1.** Summary of the data used for analysis

Region	Min. vintage	Min. date	Number of auction lots	Number of wine-vintages
Australia	1945	2001–10	663,456	32,902
Bordeaux	1945	2002–03	236,534	15,221
Burgundy	1945	2002–03	63,092	2,551
California	1978	2002–05	56,540	11,988
Italy	1955	2002–03	48,249	9,981

converted from the original auction currency to U.S. dollars using the exchange rate on the date the auction was held. Only standard 750 ml bottles were retained for the analysis.

#### IV. Analysis methods

Since the primary goal of this analysis is to distinguish market trends from price appreciation due to wine maturation, APC models are a natural approach. APC models were designed to decompose cohort (or in this case, wine-vintage) performance data into three constituent functions: a lifecycle function  $F(a)$  of age  $a$  of the wine-vintage, a quality adjustment function  $G(v)$  of vintage  $v$  (the year recorded as vintage on the bottle being sold), and an environment (market) function  $H(t)$  of calendar date  $t$ . In the case of modeling auction prices  $p(i, a, v, t)$  for label  $i$ , this can be expressed as

$$\log(p(i, a, v, t)) = F(a) + G(v) + H(t),$$

where  $G(v)$  is an overall quality adjustment relative to the vintage year. The analysis can be further segmented as  $G(i, v)$ , a quality adjustment for a specific label, if enough data exists for each label  $i$ .

APC models are well-suited to the current analysis because they emphasize estimating the nonlinear structure of the functions of age and calendar date, which is key to understanding the long-term potential for appreciation in wine prices. Linear or low-order polynomial functions of age do not provide the resolution necessary to draw meaningful conclusions about lifecycle and market impacts upon price appreciation.

The quality adjustments for a vintage year or a specific label are assumed to be constant adjustments relative to lifecycle and market. As such, a detailed analysis of the explanatory factors behind these adjustments, although interesting, is not relevant for the current analysis.

The estimation can be performed with spline approximations of the functions using maximum likelihood estimation (Mason and Fienberg, 1985; Holford, 1983, 2005) or via a Bayesian estimation of the functions represented nonparametrically with one parameter at each point in age, vintage, and calendar date (Schimid and Held, 2007). A Bayesian estimation approach was employed here because it avoids the premature smoothing of spline estimation and polynomial runaway in the tails of the functions where the data becomes thin. The use of APC models for analyzing wine auction prices has been previously described by Breeden and Liang (2017).

APC models may be understood as a specific kind of panel regression, where wine-vintages are observed repeatedly via auction lot prices. APC models can also be described within the context of hedonic regression (Rosen, 1974). The APC modeling technique is also similar to Repeat Sales Regression (Bailey, Muth, and Nourse, 1963), although APC in this research is being applied to analyzing prices directly rather than price changes because the data density is very nonuniform with time.

APC analysis of wine data is most similar to the research of Dimson, Rousseau, and Spaenjers (2015), who examine the impact of aging on wine prices and by

using a value-weighted arithmetic repeat-sales regression over 1900–2012, examining the long-term return for vintages of Haut-Brion, Lafite-Rothschild, Latour, Margaux, and Mouton-Rothschild. They looked for price appreciation due to lifecycle with wine age, environmental impacts with a calendar date, and included production and quality as explanatory variables to explain vintage variation.

We begin by applying APC models to auction price data by region to show overall appreciation for age and market changes. Previous research (Breedon, Leonova, and Bellotti, 2021) has shown that data with age, vintage, and time effects cannot be modeled without either a linear function ambiguity as with APC models or a multicollinearity problem as with regression and Cox PH models. For this analysis, the vintage function was assumed to have no net linear component, which implies that all long-term price appreciation will appear either in the market function or lifecycle function.

To explore segmented analysis, the APC models are applied to the trade volume for each label to measure the expected frequency of trades versus age and normalized for the overall data volume by calendar date. This allows us to predict the lifetime trading popularity of a wine very early in its age. The wines in each region are then segmented based upon historic and predicted popularity to re-estimate price appreciation via a popularity-segmented APC analysis within each region.

### A. Results without segmentation


The first step in the analysis was to create a baseline for comparison by performing APC analysis on all available auction data within each region for labels with enough trades for analysis. Figure 1 shows the resulting price appreciation lifecycles versus the age of wine in years. The Bayesian APC approach was employed, so the results have no smoothing. For use in forecasting for portfolio management, one would prefer to smooth these functions, guided by the confidence intervals shown as shaded regions.

The wines of most regions show an initial drop in price followed by a long-term appreciation over the subsequent 20 to 30 years before the prices flatten out. The results are shown on a log scale to make the dynamics more visible.

The market trends from the APC analysis by region are shown in Figure 2 with January 2010 normalized to 100. The original environmental functions  $H(t)$  were measured on a log scale, so the market index is computed as  $\exp(H(t))$ . The confidence intervals were not plotted in order to simplify the visual presentation, but they can be inferred from the variation in the graphs.

Some features of the graphs agree with intuition. The “Lafite Bubble” of 2011–2012 really affected all Bordeaux wines and can be seen here as a dramatic boom-bust cycle. Throughout 2017–2019, a dramatic increase in Burgundy is apparent, but without a market collapse so far. Australian wines have been in a trading range since 2012, which can nevertheless be profitable for active investors. Italian wines have seen fewer significant booms and busts, while Californian wines have been much less volatile than the others, except for a period of appreciation beginning in 2018.

Analysis of many candidate segmentations shows that the market estimates are very robust within a region. If we just measure the market appreciation across the




**Figure 1.** Price appreciation lifecycles from APC analysis, showing the average  $\log_{10}(\text{price})$  for red wine versus age for each region.

full 14-year period, the result can be susceptible to recent increases, such as with Burgundy and California. As an alternative, we computed the mean one-year market return and confidence interval, to show how volatile the market can be. A buy-and-hold strategy in this kind of market could be problematic, but volatility creates opportunities for an active trader.

Table 2 shows the market returns from January 2005 through December 2018, a 14-year span, as well as the average one-year return starting from any month in that series, and the 95% confidence interval about the mean. All of the returns are statistically significant over this period.

Note that compounding the annual return over 14 years would be dramatically higher than the observed returns from 2005 through 2018 because wine returns are strongly anticorrelated over 12 to 24 months, which is a reflection of the boom-bust cycles observed in Figure 2.



**Figure 2.** Market price indices from APC analysis for red wines from each region versus calendar date (monthly).

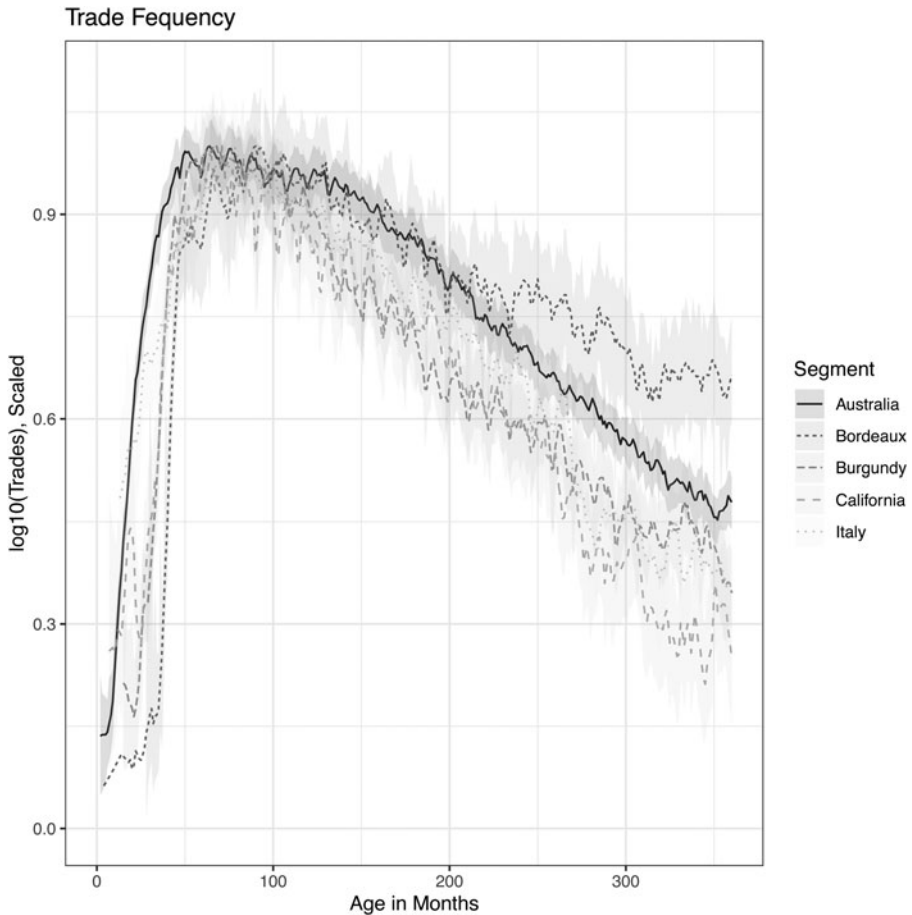
**Table 2.** Investment returns realized from the market indices in [Figure 2](#)

Region	Return (in percent)		
	2005–2018	1-year average	1-year standard error
Australia	209.1	11.0	3.3
Bordeaux	41.3	5.9	3.3
Burgundy	416.0	14.9	4.3
California	80.7	6.9	2.0
Italy	86.5	12.0	6.0

### B. Modeling wine-vintage popularity

As we explored segmenting the analysis within regions, we considered segmenting by wines with the highest initial price, from the most popular vintages and those most commonly selected in wine market baskets. All of these approaches have serious flaws. Initial price sales indicate little about long-term appreciation. Not all wines of a given vintage have the same potential for appreciation. Wines in popular market baskets are generally older and thus are not useful for understanding the appreciation potential of newer vintages.

After discussions with industry experts and dissatisfaction with the test results for the methods mentioned previously, we looked at predicting label popularity so that labels can be segmented very early in their lifecycles for the proper estimation of their appreciation value. This was done by creating an APC model for the trading



**Figure 3.** Lifecycles for trading volume versus age of a label. The original estimate was the log of trading volume, which was scaled for this graph between 0 and 1 to normalize for the differing data density between regions.



frequency of labels by estimating the trading volume versus the age of the label and trading volume by calendar date. Then, the popularity of each label,  $G_P(i, \nu)$ , is estimated as a one-parameter regression with lifecycle  $F_P(a)$  and environment  $H_P(t)$  as fixed inputs.

$$\text{Num Trades} \sim F_P(a) + H_P(t) + G_P(i, \nu)$$

In absolute terms, the lifecycle and environment reflect the density of our training data. The analysis is more insightful when measured relative to the total volume of data. Figure 3 shows these normalized lifecycles—this time with age measured monthly rather than annually, so that the variation within the year is clear. These results show that the rate of increase and decrease are quite similar across regions, except for Bordeaux, which has a significantly higher long-term trading frequency.

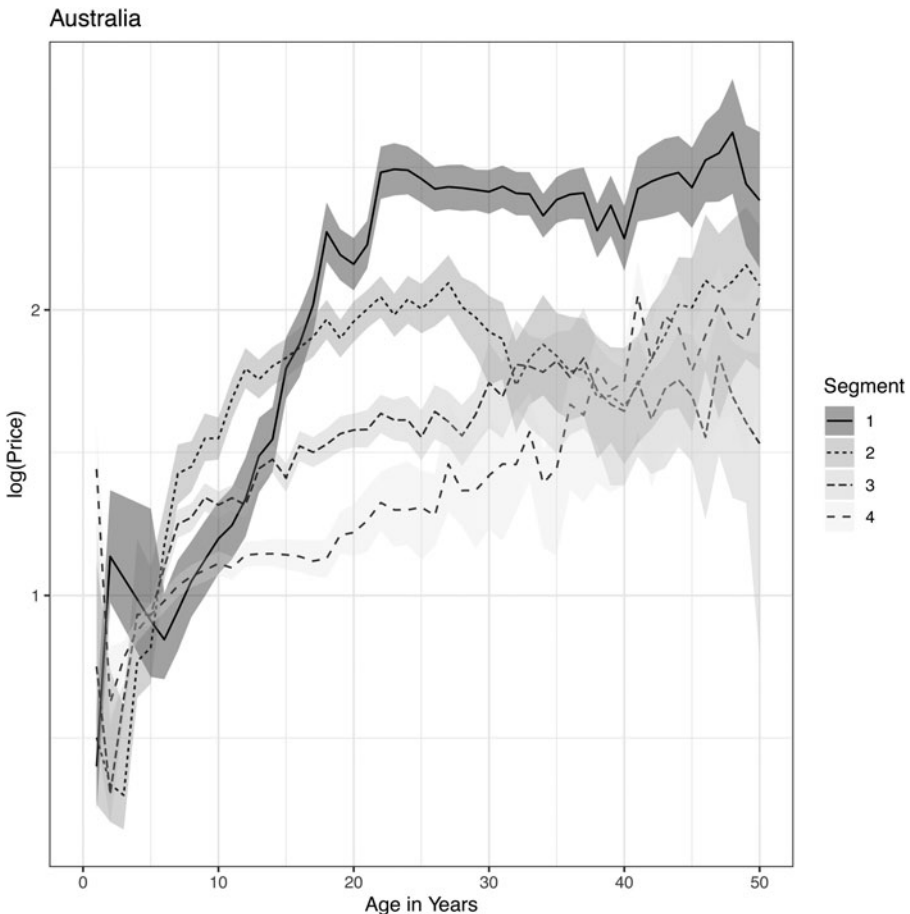


Figure 4. Price lifecycles for Australian wines as obtained from a Bayesian APC analysis applied to segmentation by popularity tiers.

### C. Results with segmentation by popularity

The value of the popularity analysis is that the label popularity  $G_P(i, v)$  can be estimated for any age and allows us to create segments with wines of different ages. Within each region, we created four segments based upon ranking the popularity: 1–25, 26–100, 101–500, and 501+. With these segments, the APC analysis of log (price) was rerun for each region, as shown in Figure 4 through Figure 8.

To glean insights from these results, compare the segmented region analysis to the unsegmented result in Figure 1. Australia is an anomaly compared to the other regions. Starting with Bordeaux, because those wines are the most popular for investors, the top 25 wines immediately begin with high prices. Prices peak around age 12 and then are largely flat thereafter. The second and third-tier wines rapidly increase in price over the first ten years and then flatten out. This analysis reveals that the perceived increase in price with the age of the vintage in Figure 1 is due to a selection

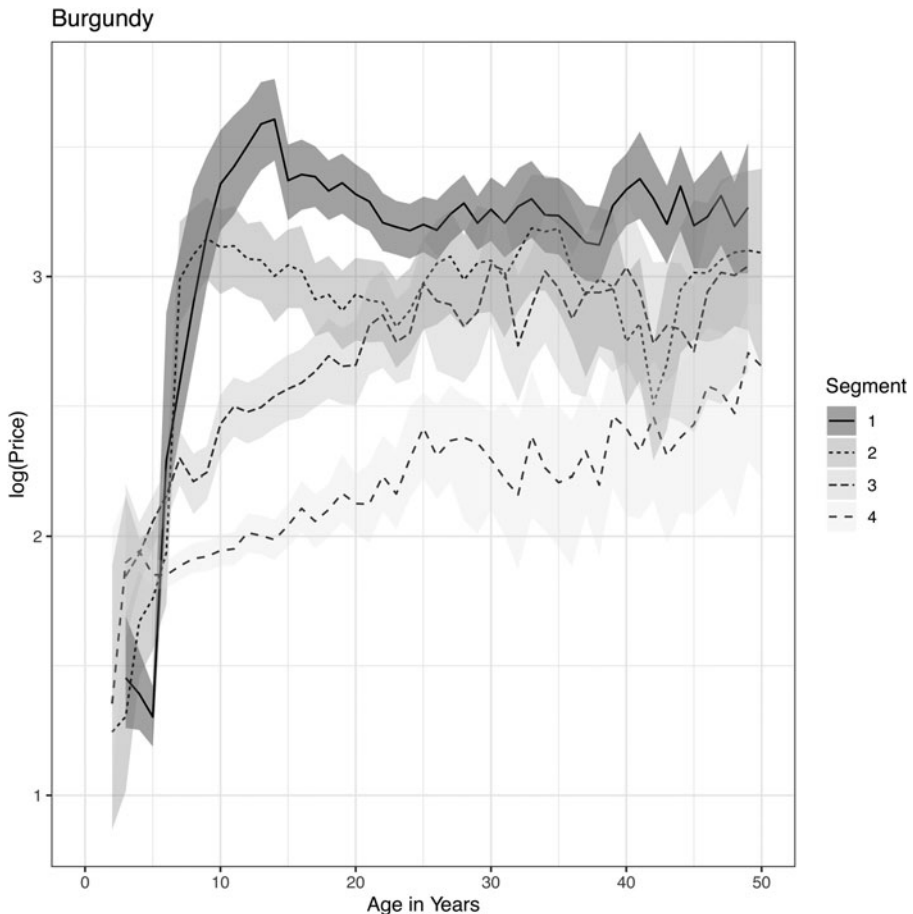




Figure 5. Price lifecycles for Bordeaux wines as obtained from a Bayesian APC analysis applied to segmentation by popularity tiers.



**Figure 6.** Price lifecycles for Burgundy wines as obtained from a Bayesian APC analysis applied to segmentation by popularity tiers.

bias. As wines aged, the higher-priced wines continued to be traded, whereas the less expensive wines did not. Therefore, the buy-and-hold strategy does not guarantee profits for Bordeaux wines and has been largely just an illusion due to selection bias.


The same effect can be seen in Burgundy wines. Tier 1 and 2 both increase rapidly in price for the first 10 years and then flatten out, considering the confidence intervals. Tiers 3 and 4 show a continued increase, but one must question whether this is due to a selection bias again. We cannot further segment the analysis for the less popular wines because the data is quickly exhausted. However, since tiers 1 and 2 are the primary wines for trading, the result is the same as for Bordeaux. Namely, price appreciation with age is not reliable, and the observed increases are better explained as an overall increase in the market for Bordeaux and Burgundy wines. The common perception was that appreciation with age was more reliable than market increases because of past boom-bust events in the market. This analysis shows no reliable increases with age, and investors must assume that all gains are susceptible to future market crashes.



**Figure 7.** Price lifecycles for California wines as obtained from a Bayesian APC analysis applied to segmentation by popularity tiers.

Californian and Italian wines largely continue the pattern just observed. Prices in these markets may peak even earlier, such as at 8 to 10 years of age, and flatten from there. The data in both markets become sparser and more volatile after 20 years, so it may be prudent to assume that prices are mostly flat thereafter.

Lastly, returning to Australia, we see a fundamentally different pattern. Only in this region do prices continue to rise for 20 to 30 years. Still, the top 25 wines plateau first, but the rise is longer than in other regions. Although Wood and Anderson (2006) also analyzed Australian wine price appreciation, their analysis focused primarily on growth conditions such as rainfall, temperature, sun, etc. Their predictive factors relative to age are considered only in linear, squared, and cubed terms. Such a low-order polynomial representation will not capture the nonlinear structure being studied here and will be unstable in the high-age tails of the function. As such, the current results reveal a structure not previously observed.



**Figure 8.** Price lifecycles for Italian wines as obtained from a Bayesian APC analysis applied to segmentation by popularity tiers.

## V. Conclusion

Assessing investment returns for wine is challenging because of the complex dynamics of wine prices and the numerous selection biases in the available data. Simple average wine indices can be deeply misleading, as recognized by previous researchers who have sought to tease apart the various drivers from auction results.

The present analysis provides a study of how selection biases may be driving spurious results in assessing wine investment returns. In addition, we found that segmenting by wine popularity is helpful in stabilizing the analysis.

**Acknowledgments.** The author would like to thank Karl Storchmann (the editor) and an anonymous reviewer for their helpful comments and encouragement.

## References

- Ashenfelter, O. (2010). Predicting the quality and prices of Bordeaux wine. *Journal of Wine Economics*, 5(1), 40–52.
- Bailey, M. J., Muth, R. F., and Nourse, H. O. (1963). A regression method for real estate price index construction. *Journal of the American Statistical Association*, 58(304), 933–942.
- Balan, T. A., and Putter, H. (2020). A tutorial on frailty models. *Statistical Methods in Medical Research*, 29(11), 3424–3454.
- Bouri, E., Gupta, R., Wong, W. K., and Zhu, Z. (2018). Is wine a good choice for investment? *Pacific-Basin Finance Journal*, 51, 171–183.
- Breeden, J. L. (2010). *Reinventing Retail Lending Analytics: Forecasting, Stress Testing, Capital and Scoring for a World of Crises*. London: Risk Books.
- Breeden, J. L., Leonova, E., and Bellotti, A. (2021). Instabilities using Cox PH models on data from systems in more than two dimensions. Unpublished manuscript, August. Download at [https://www.researchgate.net/publication/356082846\\_Instabilities\\_Using\\_Cox\\_PH\\_Models\\_on\\_Data\\_from\\_Systems\\_in\\_More\\_than\\_Two\\_Dimensions/link/618b2ad007be5f31b76223c7/download](https://www.researchgate.net/publication/356082846_Instabilities_Using_Cox_PH_Models_on_Data_from_Systems_in_More_than_Two_Dimensions/link/618b2ad007be5f31b76223c7/download).
- Breeden, J. L., and Liang, S. (2017). Auction-price dynamics for fine wines from age-period-cohort models. *Journal of Wine Economics*, 12(2), 173–202.
- Burton, B. J., and Jacobsen, J. P. (2001). The rate of return on investment in wine. *Economic Inquiry*, 39(3), 337–350.
- Chu, P. K. K. (2014). Study on the diversification ability of fine wine investment. *Journal of Investing*, 23(1), 123–139.
- Dimson, E., Rousseau, P. L., and Spaenjers, C. (2015). The price of wine. *Journal of Financial Economics*, 118(2), 431–449.
- Fogarty, J. J. (2006). The return to Australian fine wine. *European Review of Agricultural Economics*, 33(4), 542–561.
- Fogarty, J. J. (2010). Wine investment and portfolio diversification gains. *Journal of Wine Economics*, 5(1), 119–131.
- Fogarty, J. J., and Sadler, R. (2014). To save or savor: A review of approaches for measuring wine as an investment. *Journal of Wine Economics*, 9(3), 225–248.
- Holford, T. (1983). The estimation of age, period, and cohort effects for vital rates. *Biometrics*, 39(2), 311–324.
- Holford, T. (2005). Age-period-cohort analysis. In B. Everitt and D. Howell (eds.), *Encyclopedia of Statistics in Behavioral Science*. Hoboken, NJ: John Wiley & Sons, Ltd.
- Jaeger, E. (1981). To save or savor: The rate of return to storing wine. *Journal of Political Economy*, 89(3), 584–592.
- Jones, G. V., and Storchmann, K. (2001). Wine market prices and investment under uncertainty: An econometric model for Bordeaux crus class. *Agricultural Economics*, 26(2), 115–133.
- Jovanovic, B. (2008). Bubbles in prices of exhaustible resources. American Association of Wine Economists, Working Paper No. 32, December. Available from [https://wine-economics.org/wp-content/uploads/2012/10/AAWE\\_WP32.pdf](https://wine-economics.org/wp-content/uploads/2012/10/AAWE_WP32.pdf).
- Krasker, W. S. (1979). The rate of return to storing wines. *Journal of Political Economy*, 87(6), 1363–1367.
- Mason, W. M., and Fienberg, S. (1985). *Cohort Analysis in Social Research: Beyond the Identification Problem*. New York, NY: Springer.
- Masset, P., and Henderson, C. (2010). Wine as an alternative asset class. *Journal of Wine Economics*, 5(1), 87–118.
- Masset, P., and Weisskopf, J. P. (2018). Raise your glass: Wine investment and the financial crisis. In O. Ashenfelter, O. Gergaud, K. Storchmann, and W. Ziemba (eds.), *World Scientific Reference on Handbook of the Economics of Wine: Volume 1: Prices, Finance, and Expert Opinion*, 271–295. Tuck Link, Singapore: World Scientific Publishing Co Pte Ltd.
- Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34–55.
- Sanning, L. W., Shaffer, S., and Sharratt, J. (2008). Bordeaux wine as a financial investment. *Journal of Wine Economics*, 3(1), 51–71.
- Schmid, V. J., and Held, L. (2007). Bayesian age-period-cohort modeling and prediction - BAMP. *Journal of Statistical Software*, 21(8), doi: [10.18637/jss.v021.i08](https://doi.org/10.18637/jss.v021.i08).
- Storchmann, K. (2012). Wine economics. *Journal of Wine Economics*, 7(1), 1–33.
- Wood, D., and Anderson, K. (2006). What determines the future value of an icon wine? New evidence from Australia. *Journal of Wine Economics*, 1(2), 141–161.