Effects of Harvest Season Temperatures on Retail Markup of Bordeaux Wines

Brady Gartman and Erika Haguette
NYU Economics

I. INTRODUCTION AND MOTIVATION

This paper aims to provide actionable information to directly benefit wine retailers, leading to profit-increasing decision making. Our analysis goes beyond existing work by examining the relationship between weather and profits as well as the current allocation of these profits. We provide information and tools which can be used by wine retailers in the design of their strategy.

We seek to reveal whether it is always more profitable for retailers to sell the most expensive wines from the best years. Our results contradict this hypothesis, suggesting that retailers will increase profits per bottle sold by selling less expensive wines from hot years and more expensive wines from cool years, all other factors held constant. We also consider the possibility that wine retailers may be receiving some indirect benefit from selling the most expensive wines from the highest quality years.

From the publication of the Bordeaux Equation by Princeton economist Orley Ashenfelter in 1995 to the establishment of the peer-reviewed *Journal of Wine Economics* in 2006, empirical research on wine economics has gained increased exposure in recent years. Simultaneously, consumption of wine has seen marked growth, with a 108% increase in U.S. consumption alone between 1993 and 2018, from 370 million to 770 million gallons according to Rob McMillan, EVP and founder of the Silicon Valley Bank Wine Division.

As described by leading wine economist Karl Storchmann (2012), fine wine differentiates itself from other agricultural commodities and beverages in four main ways:

1. Prices regularly exceed several thousand dollars per unit;
2. it can be stored for substantial lengths of time (up to several decades for some varieties) and can increase in value over time;
3. its quality (and therefore price) are extremely sensitive to weather fluctuations; and
4. it is difficult to ascertain its quality before consumption, meaning consumers tend to rely on expert ratings.

Due to the economic nature of fine wines in general as well as data availability, we focus this research on fine wine from the Bordeaux region of southern France. Henceforth, we use “wine” to refer to Bordeaux wines exclusively.

Because of fine wine’s sensitivity to weather fluctuations, industry stakeholders have an interest in the predicted effects of climate change. According to data from the National Oceanic and Atmospheric Administration, the six hottest years ever recorded have happened since 2010. The year 2017 was the third warmest year, behind 2016 and 2015, and approximately 0.8 degrees Celsius hotter than the twentieth-century average. In fact, every year since 1977 has been hotter than that average.

For the Bordeaux region, the Intergovernmental Panel on Climate Change forecasts temperature increases between two and four degrees Celsius by the year 2050. Simultaneously, Van Leeuwen and Darriet (2016) find that, out of the twenty driest years in the region since 1955, ten were between the years 2000 and 2012. For Bordeaux wines, drier and hotter years are nearly synonymous with higher quality. Given the trend towards hotter and drier years, it is increasingly relevant that wine makers and retailers understand the effects of temperature and precipitation on their business.

II. DATA

We constructed our own cross-sectional dataset containing 729 observations of five Tier I and twelve Tier II chateaux from 106 NY retailers. When we partitioned the data at the median price (explained in Section II of this paper), the lower-priced sample had 356 observations while the higher-priced sample had 373 observations. Specific vintages included range from 1982 to 2015 based on data availability, for a total of 161 unique chateau-vintage combinations. Summary statistics for each variable used in the model are presented in Table 1.

Table 1: Variables

	Format	Mean	Range	Source
Dependent				
Markup	Continuous, non-negative	0.200	[0, 0.666]	1
Independent				
Average Temperature	Continuous (°C), Aug-Sept	32.726	[29.83, 35.11]	2
Total Average Precipitation	Continuous (cm), Aug-Sept	2.863	[1.04, 6.65]	2
Control				
Age	Integer	9.47	[3, 36]	1
Retailer	Dummy	n/a	n/a	1
Wholesale price	Continuous (US$)	380.08	[46, 2500]	3
Wine Name	Dummy	n/a	n/a	1
*Quality Score	Continuous	94.46	[86, 100]	4
*Retailer Zip	Integer	n/a	n/a	1

1. *Winesearcher* online wine database, www.wine-searcher.com
2. KMNI climate explorer, climexp.knmi.nl
3. NY State Liquor Authority, sla.ny.gov
4. Robert Parker (*Wine Advocate*), www.robertparker.com

Asterisk (*) indicates variable considered in analysis but omitted from final model

For retail prices, we enlisted the help of a software engineer to mine the leading wine search engine, *WineSearcher*. We aggregated the October 2018 retail prices of individual, standard sized (750 mL) bottles of wine for sale in New York state, excluding auction and bulk prices, as well as each retailer's name and zip code. The wines included in this dataset represent all but two chateaux classified as Tiers I or II per the industry-standard Médoc Classification of 1855 (Peppercorn 2003).

Weather data, consisting of daily temperature and precipitation measurements, is from the Royal Netherlands Meteorological Institute (KNMI) website. This organization reports official values from a single government-operated weather station in the Bordeaux-Merignac region. We converted daily readings into average monthly values for both temperature and precipitation. We report temperature as the average of the months of August and September, defined as the harvest season, because it is the time in which weather has the greatest impact on the quality of wine grapes (Storchmann 2012). Precipitation is likewise the total precipitation during the harvest season.

The New York State Liquor Authority reports wholesale prices, which by law must be consistent across retailers within the state. We calculated wine age by subtracting the vintage from the year of data collection, 2018. Finally, quality score refers to the ratings given to each wine by *The Wine Advocate*, a well-respected publication founded by master sommelier and industry expert Robert Parker.

The dependent variable of interest is percent markup, which we refer to as simply markup. This variable is a function of retail and wholesale prices, computed as

$$\text{percent markup} = \frac{\text{retail price - wholesale price}}{\text{retail price}} \qquad \text{Equation 1}$$

We converted this number into signed deviation of markup from its mean by subtracting the sample average from each observation. This does not impact the coefficient values or interpretation of results, but it does improve the communicability of these results to our audience, the retailer. Note that, although markup is dependent on wholesale price, it is linearly independent and therefore wholesale price is a valid control variable to regress against markup. Markup is an adequate proxy for profits in this analysis because it is reasonable to assume that the only cost change between stocking bottle A or bottle B is the difference in their respective wholesale prices.

We considered including quality score as a control variable. However, Ashenfelter and Jones (2012) found not only that weather is a better predictor of quality than scores, but also that, when added to the weather equation, scores become insignificant and inconsistent and do not provide any private information. Another possible control was retailer zip code, but it proved to be inferior to using the retailer name as a control because it did not add any information and increased standard errors via higher multicollinearity.

III. METHODOLOGY

To capture the effect of temperature on wine markup, we selected a standard econometric linear model. This model is commonly implemented for its simplicity of interpretation, which facilitates inferencing. The cost of using a linear model is the potential predictive accuracy of more complex functional forms, which is an acceptable tradeoff for this analysis because we favor the interpretability of linear coefficients as estimates of the marginal effect of independent variables.

In each of the models described below, we seek to estimate the marginal effect of weather on the dependent variable. The general form of such a model is given by

$$y = X\beta + Z\alpha + \varepsilon \qquad \text{Equation 2}$$

where y is the dependent variable, X is a matrix containing the independent variables (in this case temperature and precipitation), and Z is a matrix of controls. β and α represent vectors of coefficients to be estimated by ordinary least squares, and ε represents the error unexplained by the linear effects of X and Z together.

We also corrected for heteroskedastic variables in each of our models. Keeping with econometric standards, we confirmed the suspected heteroskedasticity with the White test and corrected for it using the popular HC0 correction. The technical details of this correction are beyond the scope of this research; see Long and Ervin (2000).

Ashenfelter (2008) established the relationship between harvest season weather and wine price. We began our analysis by confirming this relationship within our own dataset. These preliminary models examine the marginal effect of precipitation and rainfall during the harvest period on both wholesale and retail prices. The results of these regressions are presented in Table 2.

Table 2: Weather-Price Models

			Dependent variable	
			Wholesale Price	Retail Price
Independent variable	Average Temperature	Coefficient	0.0638	0.0432
		Std error	0.012	0.012
		p-value	0.000	0.000
	Total Average Precipitation	Coefficient	-0.1175	-0.1161
		Std error	0.008	0.008
		p-value	0.000	0.000
		R-squared	0.930	0.932

Controls for age, retailer, and wine included; SE are heteroskedasticity robust with HC0 correction

As expected, the effect of temperature on prices is positive and the effect of precipitation on prices is negative. Both are statistically significant even at significance levels of 99.9%. This result supports the consensus that hotter, drier years will yield more expensive wines, almost certainly because they tend to produce higher quality grapes. Furthermore, this result validates the legitimacy of our data as a representative sample of Tier I and II Bordeaux wines.

We next examined the relationship between markup and weather. This model is formulated similarly to the price models, although here we have substituted markup as the dependent variable. On a theoretical basis, one would expect higher-price wines of the same chateau to be more profitable, as production costs should be unaffected by the quality of the wine year-to-year. The allocation of these profits to producers, retailers, and along the supply chain between them is poorly understood.

Intuitively, one might predict that retailers appropriate at least some of these profits. This hypothesis would be confirmed by a positive relationship between temperature and markup (and likewise a negative one between precipitation and markup). Others might reasonably expect that there is no relationship between weather and markup, as non-retail firms higher up the supply chain may absorb all additional profits before retailers purchase these better wines. Our results are shown in Table 3.

Table 3: Weather-Markup Model

			Dependent variable is Markup
Independent variable	Average Temperature	Coefficient	-0.0153
		Std error	0.004
		p-value	0.000
	Total Average Precipitation	Coefficient	-0.0021
		Std error	0.003
		p-value	0.416
		R-squared	0.930

Controls for age, wholesale price, retailer, and wine included;
SE are heteroskedasticity robust with HC0 correction

These results support a rejection of both intuitive hypotheses proposed above. While the effect of precipitation on markup is not statistically significant at any reasonable significance level, the effect of temperature is significant even at the 99.9% level. Most notably, the effect of temperature on markup is *negative*. We estimate that a one degree Celsius increase in the average temperature during the months of August and September *decreases* the retail percent markup by 0.0153 points.

Given this result, we proceeded with the final model, which examines how the effect of temperature on markup varies for high- and low-priced wines. Note that at this point we focus on the effect of temperature and re-designate precipitation as a control because (i) the effect of precipitation on markup is statistically zero and (ii) looking forward, contemporary climate science predicts that average temperatures in the Bordeaux region will rise as a result of climate change. For this reason, stakeholders in the wine industry should expect the total effect of temperature on Bordeaux prices to increase over time.

We divided the sample approximately in half at the median wholesale price to distinguish the difference in the effect of temperature on markup between more and less expensive wines. We then evaluated the same model as before (with markup as the dependent variable) separately on each segment, producing the results presented in Table 4.

IV. RESULTS & PERFORMANCE

When we partition the dataset on the median wholesale price and evaluate each regression separately, the distinct effects of temperature on the markup of highly priced wines becomes clearer. These results are presented in Table 4.

Table 4: Weather-Markup Model with Partitioned Sample

		Dependent variable is Markup	
		Price < 245	Price ≥ 245
Independent variable is Average Temperature	Coefficient	-0.0072	-0.0164
	Std error	0.006	0.006
	p-value	0.210	0.004
	R-squared	0.641	0.470

Controls for age, precipitation, wholesale price, retailer, and wine included;
SE are heteroskedasticity robust with HC0 correction

The low-price data segment has no significant relationship with temperature at the standard 95% significance level. A p-value of 0.210 is sufficiently low that we interpret this relationship as statistically indistinguishable from zero. This indicates that, for the entire sample, most of the effect of temperature can be attributed to the high-price data segment. When we isolate this price range, the effect of temperature is significant at the 99.5% level. Further, the coefficient on temperature is larger in magnitude by approximately 7%.

The explanatory power of the low-price model (0.641) is stronger than the high price model (0.470), but both improve upon the whole-sample model, supporting the logic of dividing the sample in this way. Although the R-squared values of these models may seem low, this research is but one of many tools which must be used concurrently in the design of an effective retail strategy. By this reasoning, we are satisfied with our model's contribution towards increasing the breadth of actionable information available to the retailer in their decision-making process.

The coefficient on the high-price data segment indicates that a one degree Celsius increase in average August-September temperatures would decrease markups of these wines by 0.0164 percentage points. Current predictions of the climate change impacts on the Bordeaux region predict an increase in the average growing season temperature between two and four degrees Celsius by 2050. Such a systematic change, to say nothing of annual fluctuations, stands to significantly impact the profits of retailers selling these wines.

These results lead to the following conclusions and actionable insights:

1. When the Bordeaux region experiences above average harvest season temperatures (hotter years), the quality of the wine grapes is higher than cooler years. As expected, these grapes yield more expensive wines than those from cooler years, even for wines of the same variety from the same vineyard.
2. Considering that packaging and shipping costs for any particular chateau's wine should not vary significantly between vintages, this increase in consumer willingness to pay for wines from hotter years should directly increase the profitability of these wines.
3. The exact allocation of the additional profits seen in hotter years remains unknown. They are assumed to be appropriated by some firm(s) higher in the supply chain, between wine production and retail consumption. We now know, however, that low-priced wines have no residual profits

from the higher temperatures remaining for retailers. For high-priced wines, we now know that warmer years generate *lower* profits than cooler years.

4. Retailers are effectively paying a premium for selling the highest quality wines. We can think of this premium as a sort of 'luxury tax,' which is appropriated somewhere higher up the supply chain. This contradicts the intuition that the best wines to sell at retail are those from the best years.

On one hand, the most immediate explanation for this effect is that retailers relatively overvalue the benefits of stocking the most prestigious wines. We believe retailers are losing potential profits by being unaware of this overvaluation. On the other hand, higher prices can be a signal of more retailer prestige in some markets. It may be prudent for retailers to continue to invest in high priced wines if stocking these wines facilitates their marketing strategies.

In summary, this paper finds a small but significant negative effect of temperature on the retail markup of high-priced Tier I and II Bordeaux wines and no significant effect for low-priced wines. We also conclude that there is no significant effect of precipitation on markup, regardless of price.