

Chapter 25

Do More Expensive Wines Taste Better? Evidence from a Large Sample of Blind Tastings*

Robin Goldstein^{†,||}, Johan Almenberg[‡], Anna Dreber[§],
John W. Emerson[¶], Alexis Herschkowitsch[†],
and Jacob Katz[†]

[†]*Fearless Critic Media, 2011 B Bouldin Avenue,
Austin, TX 78704, USA
|| robin@fearlesscritic.com*

[‡]*Department of Economics,
Stockholm School of Economics, and the
Research Institute of Industrial Economics,
Sveriges Riksbank, SE-103 37 Stockholm, Sweden
johan.almenberg@hhs.se*

[§]*Department of Economics,
Stockholm School of Economics, and
Program for Evolutionary Dynamics,
P.O. Box 6501, 113 83 Stockholm, Sweden
anna.dreber@hhs.se*

[¶]*Department of Statistics, Yale University, USA
john.emerson@yale.edu*

Individuals who are unaware of the price do not derive more enjoyment from more expensive wine. In a sample of more than 6,000 blind tastings, we find that the correlation between price and overall rating is small and negative, suggesting that individuals on average enjoy more expensive wines slightly *less*. For individuals with wine training, however, we find indications of a non-negative relationship between price and enjoyment. Our results are robust to the inclusion of individual fixed effects, and are not driven by outliers: when omitting the top and bottom deciles of the price distribution, our qualitative results are strengthened, and the statistical significance is improved further. These findings suggest that non-expert wine consumers should not anticipate greater enjoyment of the intrinsic qualities of a wine simply because it is expensive or is appreciated by experts.

*This article originally appeared in *Journal of Wine Economics* (2008), 3(1), 1–9.
©The American Association of Wine Economists, 2008.

1. Introduction

When symbolic content is an important part of consumption, the enjoyment of a good might become decoupled from its innate qualities. The symbolic content of a price tag has been emphasized in marketing research (e.g., Cialdini, 1998). At the same time, when goods with similar characteristics differ in price, a reasonable prior is that the more expensive good will, on average, be of a higher quality. People have been shown to expect a positive correlation between price and quality (e.g., Rao and Monroe, 1989). Consistent with this expectation, a meta-analysis reports positive correlations between price and quality ratings for most, but not all, of 1,200 product markets, but also finds that the range of these correlations is very large (Tellis and Wernerfelt, 1987).

For some goods, there is much heterogeneity in consumer tastes, making it harder to infer quality from revealed preferences. Nonetheless, a reasonable prior is that consumers on average will derive more enjoyment from the good with the higher price. Wine seems to be a good where consumer tastes are highly heterogeneous (Amerine and Roessler, 1976; Lecocq and Visser, 2006). While individuals may frequently disagree over which wine they prefer, the above hypothesis suggests a positive correlation between the enjoyment of a wine and its price.

A number of studies have reported positive correlations between price and subjective appreciation of a wine for wine experts (e.g., Oczkowski, 1994; Landon and Smith, 1997; Benjamin and Podolny, 1999; Schamel and Anderson, 2003; Lecocq and Visser, 2006). Non-experts, however, may not be particularly sensitive to some of the refinements that are held in high esteem by wine aficionados. Weil (2001, 2005) uses the following experimental setup: two bottles of wine are poured into four containers. Tasters are then given three of the containers and asked to distinguish which one differs from the other two. A random guess has 1/3 chance of being correct. In Weil (2001), the two wines are identical apart from year, but one wine is from a “good” vintage, and the other from a “bad” vintage.¹ The tasters get it right 41% of the time — only marginally better than a random guess.² In Weil (2005), the wines are a reserve bottling and a regular bottling, from the same producer and year. The fraction of correct answers is merely 40%. Moreover, Weil

¹Weil uses pairs for which the famous wine critic Robert Parker has rated one of the bottles “average” to “appalling” and the other bottle “excellent” to “the finest.”

²All of the significant difference is driven by the testers’ ability to distinguish between the good and bad vintages from Bordeaux Pomerol.

finds that even when tasters can distinguish between the vintages, they are about as likely to prefer the good one as the bad one. And among those who can distinguish the reserve bottling from the regular bottling, only half prefer the reserve. In both cases, the wines differ in price by an order of magnitude. These experiments highlight the discrepancy between experts and non-experts and the subjectivity of the wine experience.

Extrinsic factors, such as peer consumption and marketing actions, can also influence how a good is experienced. The price tag may in itself be such a factor. Recent research has shown that individuals appreciate the same wine more when they think that it is more expensive (Brochet, 2001; Plassmann *et al.*, 2008). In other words, the price of a good affects the experienced utility derived from that good. Thus, to test the conjecture that the price of wine and the enjoyment of its intrinsic qualities are positively correlated, we need to examine the enjoyment of wine when individuals are unaware of the price.

In this chapter, we use a large sample of more than 6,000 U.S. blind tastings conducted by food and wine critic Robin Goldstein. We investigate the relationship between price and subjective appreciation of wines when the price is unknown to the taster. Subjective appreciation is measured by overall ratings assigned to wines by individual participants. Blind tastings offer the opportunity to isolate the experience of the wine itself from psychological confounds related to its price, presentation or published expert ratings.

Our main finding is that individuals who are unaware of the price do not, on average, derive more enjoyment from more expensive wine. In fact, unless they are experts, they enjoy more expensive wines slightly *less*. Our results are robust to the inclusion of individual fixed effects, and are not driven by outliers: when we omit the extremes of the price distribution our results are even more pronounced.

This chapter organized as follows. In Section 2, we describe our data. In Section 3, we present our econometric model and report our results. We also perform a robustness check. We conclude in Section 4, where we discuss some implications of our results and suggest directions for future research.

2. Data

The data set contains 6,175 observations from 17 blind tastings organized by Robin Goldstein. The blind tastings took place in the U.S. between April 2007 and February 2008. In total, 506 participants tasted wine flights composed from 523 different wines. The wines were presented in a double-blind manner, so that neither the person serving the wine nor the person tasting

the wine knew the identity, price, or any other characteristics of the wine aside from its color. Each taster assigned an overall rating to every wine tasted prior to discussing the wines with the rest of the group, and was not permitted to change his or her answer after discussion. The rating was the response to the question “Overall, how do you find the wine?” and the available answers were “Bad,” “Okay,” “Good,” and “Great.” In the data, these alternatives are coded from 1 to 4, with 1 corresponding to “Bad” and 4 corresponding to “Great.”³

The price per bottle ranged from \$1.65 to \$150. The prices are average retail prices and were obtained from www.wine-searcher.com. The wines represent a broad variety of types (e.g., red, white, rosé, and sparkling), country origins, and grapes.

The participants were unpaid volunteers from 21 to 88 years of age. Selection bias is a concern with any voluntary subject pool, and we have no reason to think that this is an exception. It is quite likely that the sample contains an over-representation of highly educated individuals, and an over-representation of individuals working in the food and wine industries. Nonetheless, the size of the sample and the general diversity of the tasters lead us to hope that inference will not be too restricted. For a more extensive description of the blind tastings, please see Chapters 8 and 9 and Appendix 1 in Goldstein (2008).

3. Results

Throughout the regression analysis, we use both an ordered probit estimator and a linear estimator (OLS). In both cases, we consistently use robust standard errors. The ordered probit estimator is particularly well suited to an ordinal dependent variable, but we find that OLS also performs well, and yields estimates that are easier to interpret. In any case, the two models generate highly consistent results. The dependent variable is the overall rating, measured on a scale from 1 to 4, with 4 being the highest rating. The main independent variable is the price variable, expressed as the natural logarithm of the average retail price per 750 mL of the wine in U.S. dollars.⁴

³Tasters ticked one of four boxes. In about 3% of the sample, tasters ticked in between two boxes, suggesting a rating somewhere in between the two responses. For simplicity, we dropped these observations from the regression. Including them makes no difference to our qualitative results, and the changes to the estimates are negligible.

⁴If we didn't do this, we would be expecting a one dollar increase to have the same effect at the \$5 price level as at the \$50 price level. We get similar qualitative results using the

In Model 1, we regress the overall rating assigned to wine i , by individual j , on the price of the wine. About 12% of participants had some wine training, such as a sommelier course. In Model 2, we allow for the possibility that these “experts” rate wines in a different manner. We include a dummy variable for being an expert, as well as an interaction term for price and the expert dummy. In a linear regression, this allows both the intercept and the slope coefficient to differ for experts and non-experts. In terms of the linear specification, we can write these two models as

$$Y_i = \beta_0 + \beta_1 \ln(p_i) + \epsilon_i \quad (1)$$

and

$$y_{ij} = \beta_0 + \beta_2 \ln(p_i) + \beta_2 EXPERT_j + \beta_3 \ln(p_i) * EXPERT_j + \epsilon_i \quad (2)$$

where P_i is the price of wine i , and $EXPERT_j$ is a dummy variable indicating if taster j has wine training. If individuals found that more expensive wines tasted better, the correlation between overall rating and price would be positive. In our sample, this is not the case: for both the ordered probit estimates and the OLS estimates, the coefficient on price is negative. In Model 1, the OLS coefficient is about -0.04 , implying that a 100% increase in $\ln(\text{price})$ is associated with a 0.04 reduction in the overall rating (Table 1, column (c)). The negative effect for more expensive wines is statistically significant.

For non-experts, the relationship between price and overall rating is negative; for experts, however, this is not the case. Our estimates of Model 2 show that the correlation between price and overall rating is positive — or, at any rate, non-negative — for experts (Table 1, columns (b) and (d)). The price coefficient for non-experts is still negative, of about the same size as before, and with greatly improved statistical significance. The coefficient on the $\ln(\text{price}) * \text{expert}$ interaction term is highly statistically significant (ordered probit p -value: 0.017; OLS p -value: 0.015). For experts, the net coefficient on price is the sum of the two (i.e., about 0.11 for the ordered probit and 0.09 for OLS). This net coefficient is significantly different from zero, but only at the 10% level (ordered probit p -value: 0.099; OLS p -value: 0.095), despite the large sample size.

In sum, we find a non-negative relationship between price and overall rating for experts. Due to the poor statistical significance of the price coefficient

dollar prices, but the statistical significance of the coefficients is not as good (but still significant).

Table 1. Dependent Variable: Overall Rating

	Ordered probit		OLS	
	(a)	(b)	(c)	(d)
ln(Price)	-0.047 (0.039)**	-0.061 (0.013)**	-0.038 (0.038)**	-0.048 (0.012)**
ln(Price)*Expert		0.171 (0.017)**		0.138 (0.015)**
Expert		-0.558 (0.001)***		-0.448 (0.001)***
Constant			2.297 (0.000)***	2.337 (0.000)***
<i>N</i>	5986	5972	5986	5972
<i>R</i> ² /pseudo- <i>R</i> ²	0.000	0.002	0.001	0.005

Notes: Robust *p*-values in parentheses. *significant at 10%; **significant at 5%; ***significant at 1%.

for experts, it remains an open question whether this coefficient is in fact positive.

How large are these price effects? The coefficients are of a moderate magnitude, but non-negligible, given that wine prices cover a large range — both in our sample and in the general wine market. Suppose we have two wines, A and B, and Wine A costs ten times more than Wine B in dollar terms. In terms of a 100-point scale (such as that used by *Wine Spectator*), the OLS estimation of Model 2 predicts that non-experts will assign an overall rating that is four points lower for wine A, whereas experts will assign an overall rating that is seven points higher.⁵

In addition, the coefficient on the expert dummy is negative, quite sizeable, and statistically significant (OLS expert dummy coefficient: -0.448; *p*-value: 0.001). In other words, the OLS estimation of Model 2 consists of two linear relationships, one with a higher intercept (2.337) but a negative slope (-0.048), and one with a lower intercept (1.889) but a positive slope (0.090). The point where the two lines cross each other is the price level at which experts and non-experts are expected to assign the same rating. If we take the model literally, this point occurs at the price of \$25.70 (i.e., $\ln(\text{price}) = \text{approx. } 3.25$). At this price, the model predicts that both groups

⁵If the dollar price increases by a factor of 10, $\ln(\text{price})$ increases by about 2.3. Hence the predicted effect on the overall rating of tenfold increase in the dollar price is 2.3 times the $\ln(\text{price})$ coefficient for non-experts and experts, respectively, adjusted to a 100 point scale.

Table 2. Individual Fixed Effects

	Ordered probit		OLS	
	(a)	(b)	(c)	(d)
ln(Price)	-0.070 (0.007)***	-0.089 (0.001)***	-0.050 (0.009)***	-0.064 (0.002)***
ln(Price)*Expert		0.209 (0.011)**		0.151 (0.013)**
Individual fixed effects	Yes	Yes	Yes	Yes
Constant			2.487 (0.000)***	2.183 (0.000)***
<i>N</i>	5986	5972	5986	5972
<i>R</i> ² /pseudo- <i>R</i> ²	0.080	0.081	0.181	0.182

Notes: Robust *p*-values in parentheses. *significant at 10%; **significant at 5%; ***significant at 1%.

will assign a rating of about 2.2. Below this price, the model predicts that experts will assign lower ratings to a wine than non-experts, and vice versa.

We also test a third model, including individual fixed effects. In terms of the linear specification, Model 3 can be written as

$$Y_{ij} = \beta_0 + \delta_j + \beta_1 \ln(P_i) + \beta_2 \ln(P_i) * EXPERT_j + \varepsilon_i \quad (3)$$

where δ_j is a dummy for each individual taster. Including individual fixed effects has very little effect on the qualitative results and the minor differences only serve to reinforce our earlier conclusions, as both the negative effect for non-experts and the positive effect for experts become slightly stronger. These results are presented in Table 2. For each of the four regressions in Table 2, a Wald test rejects that the fixed effects are jointly equal to zero by a wide margin (*p*-value < 0.001), suggesting that this is a suitable addition to the model.

To make sure that our results are not driven by wines at the extreme ends of the price distribution, we also run our regressions on a reduced sample, omitting the top and bottom deciles of the price distribution. Given the broad range of prices in the sample, this is an appropriate precaution. The remaining wines range in price from \$6 to \$15.

Using the reduced sample, we estimate Model 2 (Table 3, columns (a) and (c)) and Model 3 (Table 3, columns (b) and (d)). This produces consistent and even more pronounced estimates. The coefficient on price is still negative, and in each case larger than when using the full sample. The statistical significance of the coefficients improves further, and the R-squared is higher.

Table 3. Reduced Sample, With and Without Individual Fixed Effects⁶

	Ordered probit		OLS	
	(a)	(b)	(c)	(d)
ln(Price)	-0.225 (0.001)***	-0.173 (0.019)**	-0.182 (0.001)***	-0.122 (0.025)**
ln(Price)*Expert	0.523 (0.002)***	0.515 (0.006)***	0.421 (0.002)***	0.364 (0.009)***
Expert	-1.301 (0.000)***		-1.044 (0.000)***	
Individual fixed effects	No	Yes	No	Yes
Constant			2.622 (0.000)***	1.910 (0.000)***
<i>N</i>	4817	4817	4817	4817
<i>R</i> ² /pseudo- <i>R</i> ²	0.003	0.094	0.007	0.206

Notes: Robust *p*-values in parentheses. *significant 10%; **significant at 15%; ***significant at 1%.

In sum, we use the reduced sample to check the robustness of our results with regard to mid-range price levels. Based on the above, we conclude that our results are not only robust but in fact even more pronounced when omitting observations at the extremes of the price distribution.

4. Conclusions

The pleasure we get from consuming wine depends both on its intrinsic qualities such as taste and smell and external attributes such as price and presentation. One may argue that the former influences our subjective appreciation through a bottom-up process, where the sensory apparatus plays a key role, and that the latter works through a top-down process, where beliefs and expectations about quality are important determinants.⁷

In this chapter we have explored the bottom-up effects by looking at how participants in blind tastings rate wines. We find that, unless they are experts, individuals who are unaware of the price enjoy more expensive wines slightly less.

⁶The reduced sample excludes the top and bottom deciles of the price distribution.

⁷This, in turn, might depend on ulterior motives such as status concerns. Wine as a status signal, and the prospect that expensive wine could function as a positional good, is discussed in Goldstein (2008), Chapter 5.

There is a large relevant literature related to the influence of extrinsic signals on taste experience. Lee *et al.* (2006) look at how knowledge of a beer's ingredients (normal beer with added vinegar) can affect subjective appreciation. They show that the timing of the information plays a substantial role. One group of tasters is told about the vinegar, tastes the beer, and assigns ratings. A second group is told about the vinegar after tasting the beer, but before the ratings are assigned. On average, individuals in the first group assign significantly lower ratings, suggesting that informing participants about the vinegar influences the experience in itself. Using functional magnetic resonance imaging (fMRI), McClure *et al.* (2004) find that having the subject's favorite brand's name on a drink makes it taste better than if it is unlabeled. In another fMRI study, Plassmann *et al.* (2008) test whether marketing actions such as changes in the price can influence the experienced pleasantness of a product such as wine. Testers are given different wines that they are told differ in price. In reality, some of the wines are the same but simply presented with different prices. Prices are found to correlate positively with experienced pleasantness, measured through both subjective reports and tMRI scans.

Marketing provides one channel through which consumers can be influenced to buy certain wines. But it is not the only one: wine critics/experts may also play a role in affecting wine prices and shaping consumer preferences. For example, Hadj Ali *et al.* (2007) find a positive effect of wine critic Robert Parker's ratings on the price of Bordeaux wine.

There is, however, some research expressing skepticism towards wine ratings and their use for the average wine drinker. According to Quandt (2007), many wine ratings do not actually convey any information, nor is there substantial agreement in ratings by experts. Consistent with this view, Weil (2007) investigates whether wine descriptions by experts convey information to wine consumers. This is tested by having testers match wine descriptions to wines. In a similar setup to Weil (2001, 2005), tasters are asked to distinguish the odd one out of three different glasses of wine. Only about 50% of the participants in Weil (2007) can distinguish the odd one out, and of those who manage to do it, only about half can correctly match a wine critic's description of the wine with the wine itself — which is no better than a random guess.

Our results indicate another reason for why the average wine drinker may not benefit from expert wine ratings: He or she simply doesn't like the same types of wines as experts. This is consistent with Weil (2001, 2005), who finds that even among the subset of tasters who can distinguish between

good and bad vintages, or reserve or regular bottlings, they are as likely to prefer the “better” one as the “worse” one.

These findings raise an interesting question: is the difference between the ratings of experts and non-experts due to an acquired taste? Or is it due to an innate ability, which is correlated with self-selection into wine training?⁸ Investigating this further would be a fruitful avenue for future research.

In sum, in a large sample of blind tastings, we find that the correlation between price and overall rating is small and negative. Unless they are experts, individuals on average enjoy more expensive wines slightly less. Our results suggest that non-expert wine consumers should not anticipate greater enjoyment of the intrinsic qualities of a wine simply because it is expensive or is appreciated by experts.

Acknowledgments

The findings reported here are discussed at length in Robin Goldstein’s book *The Wine Trials* (Fearless Critic Media, 2008). We thank Jacopo Anselmi, Zoe Chance, Shane Frederick, Richard Friberg, Barry Goldstein, Erik Grönqvist, Daniel Horwitz, Roy Ip, Magnus Johannesson, Thomas Pfeiffer, Hal Stubbs, Sue Stubbs and an anonymous referee for helpful comments and suggestions. Johan Almenberg thanks the Ragnar and Torsten Söderberg Foundations for financial support, and Johan Almenberg and Anna Dreber thank the Jan Wallander and Tom Hedelius Foundation for financial support. The Program for Evolutionary Dynamics is sponsored by J. Epstein.

References

- Amerine, M.A., and Roessler, E.B. (1976). *Wines: Their Sensory Evaluation*. San Francisco: W.H. Freeman and Company.
- Benjamin, B.A., and Podolny, J.M. (1999). Status, quality, and social order in the California wine industry. *Administrative Science Quarterly*, 44, 563–589.
- Brochet, F. (2001). Chemical object representation in the field of consciousness. Working Paper, General Oenology Laboratory, Talence, France.
- Cialdini, R.B. (1998). *Influence: The Psychology of Persuasion*. New York: Collins.
- Goldstein, R. (2008). *The Wine Trials*. Austin: Fearless Critic Media.
- Hadj Ali, H., Lecocq, S., and Visser, M. (2007). The impact of gurus: Parker grades and *en primeur* wine prices. American Association of Wine Economists, AAWE Working Paper No. 1.

⁸For a further discussion, see Chapter 4 of Goldstein (2008).

- Landon, S., and Smith, C.E. (1997). The use of quality and reputation indicators by consumers: The case of Bordeaux wine. *Journal of Consumer Policy*, 20, 289–323.
- Lecocq, S., and Visser, M. (2006). What determines wine prices: Objective vs. sensory characteristics. *Journal of Wine Economics*, 1(1), 42–56.
- Lee, L., Frederick, S., and Ariely, D. (2006). Try it, you'll like it — the influence of expectation, consumption, and revelation on preferences for beer. *Psychological Science*, 17(12), 1054–1058.
- McClure, S.M., Li, J., Tomlin, D., Cypert, K.S., Montague, L.M., and Montague, P.R. (2004). Neural correlates of behavioral preference for culturally familiar drinks. *Neuron*, 44, 379–387.
- Oczkowski, E. (1994). A hedonic price function for Australian premium wine. *Australian Journal of Agricultural Economics*, 38, 93–110.
- Plassmann, H., O'Doherty, J., Shiv, B., and Rangel, A. (2008). Marketing actions can modulate neural representations of experienced pleasantness. *Proceedings of the National Academy of Sciences*, 105(3), 1050–1054.
- Quandt, R.E. (2007). On wine bullshit: Some new software. *Journal of Wine Economics*, 2(2), 129–135.
- Rao, A.R., and Monroe, K.B. (1989). The effect of price, brand name, and store name on buyers' perceptions of product quality: An integrative review. *Journal of Marketing Research*, 36, 351–357.
- Schamel, G., and Anderson, K. (2003). Wine quality and varietal, regional and winery reputations: Hedonic prices for Australia and New Zealand. *The Economic Record*, 79(246), 357–369.
- Tellis, G.J., and Wernerfelt, B. (1987). Competitive price and quality under asymmetric information. *Marketing Science*, 6(3), 240–253.
- Weil, R.L. (2001). Parker vs. Prial: The death of the vintage chart. *Chance*, 14(4), 27–31.
- Weil, R.L. (2005). Analysis of reserve and regular bottlings: Why pay for a difference only the critics claim to notice? *Chance*, 18(3), 9–15.
- Weil, R.L. (2007). Debunking critics' wine words: Can amateurs distinguish the smell of asphalt from the taste of cherries? *Journal of Wine Economics*, 2(2), 136–144.