How to really make a wine shine?
Pour a drop of manipulation into it!

Philippe Masset[1] & Steffen Raub[2]

November 2022

Abstract

Is it possible to exploit cognitive biases so that a non-professional taster prefers one wine to several other absolutely identical wines? To address this question, three complementary experiments were carried out. Each time, five wines were tasted blind in a tasting laboratory by 24 to 34 tasters. The results show that the participants did not notice that they were tasting the same wine. Moreover, by giving them information, not only their expectations but also their evaluations were altered. We show that with a little manipulation, it is possible to modify the ranking between different wines. It is also possible to get tasters to prefer a wine over other absolutely identical wines. Finally, a surprising finding was that experienced tasters express stronger opinions and adapt their evaluations more strongly after being given manipulative information on the wines they taste.

JEL classification: C91, Q19, Z10

Keywords: wine rating; manipulation; tasting; cognitive biases; willingness to pay

[1] EHL Hospitality Business School, HES-SO, University of Applied Sciences and Arts Western Switzerland, Route de Cojonnex 18, 1000 Lausanne, Switzerland. Email: philippe.masset@ehl.ch.

[2] EHL Hospitality Business School, HES-SO, University of Applied Sciences and Arts Western Switzerland, Route de Cojonnex 18, 1000 Lausanne, Switzerland. Email: steffen.raub@ehl.ch.

1. Introduction

Research on wine tasting, evaluation and expertise has grown significantly over the past 20 years. This is due to the fact that wine offers a fascinating laboratory for studying research questions that extend beyond this noble beverage. Wine is indeed an experience good, with a strong cultural and aesthetic dimension. It is complex from both a chemical and sensory viewpoint, and comes in a variety of forms (grape variety/blends, style and aromatic profile, etc.). Research on the topic reflects the multifaceted nature of wine as it spans fields as diverse as food and beverage sciences, economics, marketing, enology, and psychology (see Spence (2020) for a recent review). Storchmann (2012) even identifies research on wine evaluation and expertise as among the three most relevant in wine economics.

Spence (2020, p. 2) notes that "a wide variety of cognitive and perceptual factors [...] influence the wine-drinking experience." This combination between a psychological (cognitive) and a mostly experienced-based (sensory) dimension implies that wine evaluation is a difficult exercise which requires an expertise that substantially differs from other realms (Ashton, 2017). Thus, wine experts have become influential in the wine market (Masset et al. (2015), Cardebat & Livat (2016)). Yet recent studies suggest that these individuals do not necessarily have a better sensory/perceptual capacity than wine novices. They distinguish themselves primarily by their ability to express their feelings using their conceptual/semantic wine knowledge (Spence & Wang, 2019).

Rodrigues & Parr (2019) illustrate the role of the taster's cultural background in wine appreciation. Their article shows, more generally, that the quality of a wine is not absolute, it depends on a frame of reference and the way the taster reacts to a variety of stimuli. Niimi et al. (2017) show, for example, that providing a description of the wine positively affects the emotions induced by its tasting (see also Niimi et al. (2019) for a brief review of the link between wine and emotions). Parr (2019) notes that "wine is as cerebral as it is sensual". As a result, it is not surprising to see that tasters, whether novices or experts, are sensitive to a set of biases induced by the color of the wine, the weight of the bottle, or its price. Goldstein (2019) shows that providing information about price biases tasters' evaluations. He identifies an asymmetric effect, with negative information (i.e., presenting the wine as cheap, the so-called nocebo effect) affecting the expectations and therefore the evaluation of tasters more strongly than positive information (the so-called placebo effect). Goldstein's article complements the

research of Plasmann et al. (2008) who show that price affects tasters' ratings but also their pleasure (measured via activity in the cortex). Actually, even when intrinsic cues are available (through sensory experience), extrinsic cues (e.g., color, prices, or expert ratings) seem to remain the key drivers of quality evaluation (Veale & Quester, 2008). Aqueveque (2018) shows, however, that experts use extrinsic cues less extensively than novices to infer quality. He further demonstrates the presence of a Dunning-Kruger (1999) effect with less (more) competent tasters tending to overestimate (underestimate) their knowledge.

Our study aims at examining the biases illustrated in the literature and, more specifically, to test whether the expectations of tasters can be manipulated to make them like certain wines more than others. We proceed in two steps. First, we test the ability of participants to notice that several wines served to them are identical. Serving identical wines is essential for ensuring that expressed preferences cannot be justified by actual differences in quality or style between wines. Next, we examine whether it is possible to direct tasters' preferences (expressed via quality ratings and willingness-to-pay) towards particular wines by altering their expectations. To do so, we influence tasters' expectations via extrinsic cues related to price, expert rating, and prestige of the wine (Ashton, 2017).

Three experiments were conducted. It should be noted that our studies deviate somewhat from the ideal put forward by Goldstein (2019) because they rely partially on deception. Depending on the experiment, deception is, as in Goldstein (2019), implicit – in the sense that participants drink the same wine several times without suspecting it, or explicit – in the sense that they are given deliberately wrong information. In any case, each tasting contains an element of manipulation.

The results are surprisingly clear-cut. Absolutely no one identified that some of the wines were identical in any of the three experiments. Interestingly this result is consistent with the study of Goldstein (2019) in which none of the participants seem to have identified that the two wines they were served were actually the same. The first experiment demonstrates that one can easily alter participants' expectations and thereby modify their preferences. In this experiment, we first served three wines to the participants, we then shared information about the ratings of the wines already tasted with them, and we finally served two last wines without giving any further information to the participants. They did not notice that the last two wines were identical to the ones they had just before and, more strikingly, they also expressed strong preferences for these last two wines. The second experiment shows that one can manipulate the expectations of tasters by giving them wrong information. One can thereby influence tasters'

evaluations and therefore their ranking of the wines they taste. It must however be noted that giving wrong information is enough to modify expectations and alter the ranking of the wines, but it is not enough to completely erase the differences between two wines. The third experiment complements this result and reveals that, with some manipulation, one can easily orient the preferences of the tasters towards a particular wine. Our results further illustrate that those who self-evaluate as "good tasters" tend to adapt their ratings and prices more strongly to the information given to them. This can be explained in two ways: (i) they express more assertive opinions because they feel more experienced, and (ii) they may also feel the need to demonstrate that they actually know wine by aligning their opinion with what they think is correct. This suggests that one of the differences between novices and more skilled consumers is that the latter express stronger opinions.

2. Research design

This section first clarifies our research agenda and the questions we examine through the experiments. The context in which the experiments took place is then presented. Finally, the experiments are discussed in detail, making the link to the research questions.

2.1. Research agenda

As explained above, tasting and rating a wine is a technical and complicated exercise. The brain plays an important role and contributes to make the exercise difficult. Indeed, various cognitive biases have been illustrated, and even professional tasters are affected by them. Expectations play a key role. The information available to tasters contributes to their expectations and consequently has an impact on their evaluations. In this study, we address the following three research questions:

I. What is the effect of wine tasters' expectations on their evaluation (quality rating and willingness-to-pay) of a wine?

II. Is it possible to reduce differences in the evaluations of objectively different wines by manipulating wine tasters' expectations?

III. Is it possible to induce differences in the evaluation of objectively identical wines by manipulating wine tasters' expectations?

The first question is directly inspired by Siegrist & Cousin (2009). They examine and demonstrate that information on wine ratings influences the expectations of tasters and thereby their ratings and willingness-to-pay. The purpose of this question is therefore to examine

whether we obtain results consistent with those documented in the literature, before tackling more original questions.

The second question incorporates elements of manipulation. Some participants receive incorrect information about the rating or the price of the wines they taste. The objective is to see if this incorrect information can be sufficient to smooth out the differences in evaluations between different wines. Specifically, the goal is to assess if the manipulation of expectations can cause a clearly better wine to end up being rated similarly (or worse) to a lesser wine. This question extends the analysis of Plassmann et al. (2008) who show that tasters will tend to prefer a wine that is supposed to be more expensive as compared to another, if both wines are actually identical. Note that this question requires the use of deception, which is unusual in economics, but not in psychological sciences. Given that our ultimate goal is to examine whether tasters can be manipulated on purpose, it is necessary and even desirable to incorporate deception in this context.

The third question is the logical conclusion of the two previous ones. It aims at examining if by using manipulation one can predetermine the wine that tasters will evaluate best. This is a completely new question, but the experiment (see below) incorporates elements already studied, notably in Wang & Spence (2019), who examine the role of wine color by coloring white wine into rosé. All participants, from beginners to experts, had more difficulty assessing the fake rosé. But interestingly experts appeared to be more influenced by the color. The authors argue that this could be explained by the fact that experts consider the color as a piece of information about the wine. In an earlier study, Parr et al. (2003) showed that experts are better at identifying "masked" red wines than social drinkers but they nevertheless remain affected by the coloring.

For the three questions and in the three experiments, we also test for potentially different behavior between tasters who evaluate themselves as good or very good and others who consider themselves rather novices.

2.2. Context

The three experiments took place during an elective course in "Wine Economics" taught in a European hospitality business school. All the participants had already followed courses on wine knowledge, enology and wine service, and had chosen to take this elective course. This means that the profile of the participants is rather homogeneous. However, some of the participants belong to wine committees and participate more frequently in tastings. This makes

it possible to examine whether wine knowledge affects the tasters' reactions to a change in their expectations. A questionnaire during the first session of the course (week 1) is used to collect this information.

The Wine Economics course runs for 5 weeks (two 3-hour sessions per week). It typically includes two tastings in addition to the experiment. These tastings are conducted with guest speakers and focus on specific wine regions. During these tastings, participants are invited to share their appreciation of the wines through a standardized questionnaire. The same questionnaire is used in the experiments. The experiments always take place during the last session of the course (week 5). This means that the students are already familiar with the questionnaire and its content. This reduces the risk of misinterpretation.

In order to prevent students from engaging in strategic behaviors (e.g., trying to identify if there is a trick), the experiments did not take place every semester. The first one took place in 2018, the second in 2021, and the third in 2022. Moreover, in order not to give students the impression that the tasting of the semester may include tricks, it is organized in the exact same way as the other tastings that take place earlier during the course: students are asked to taste silently and respond individually, but they are allowed to leave the class if necessary, or move between tasting rounds if they wish. As explained below, this may result in attrition, but this is the price to pay if one wants to make sure that participants answer in an unbiased way.

Finally, the experiments take place in a tasting laboratory with individual booths. Overall, the conditions for tasting and data collection can be considered as excellent.

2.3. Experiments

Table 1 presents an overview of the three experiments. They share several characteristics. In particular, in all three cases, five wines were tasted. Also, in each experiment, several of the wines tasted were actually identical.

The first experiment was organized to answer the first research question. This experiment consists of tasting five wines, but only three are actually different (two wines are served twice). All wines were tasted blind. For the first three wines (round 1 - wine A, round 2 – wine B, and round 3 – wine C), no information was given to the participants. The Wine Advocate's ratings of the first three wines were provided to participants after they had been tasted. This information has, of course, an effect on tasters' expectations. The wines presented in rounds 4 and 5 (wine B and wine A) were the same as those presented in rounds 2 and 1, respectively.

This experiment does not involve any deception or manipulation. The purpose is just to examine the role of expectations on participants' ratings and willingness-to-pay.

< Insert Table 1 around here >

The second experiment examines the second research question. Five wines were served. Like in the previous experiment, two wines were served twice. But, here, a bit of deception is used. Indeed, the class is divided into two groups without this information being given to the participants. In the first round, all participants receive the same wine (A). In the second round, the first group receives wine B while the second group receives wine C. The rating of wine B is also communicated to all participants. Thus, half of the participants have correct information (those who were served wine B) and the other half have incorrect information (those wo were served wine C). The deception comes in here. In the third round, wine C is served to the first group while the other group receives wine B. The rating of wine C is communicated. Rounds 4 and 5 are similar to rounds 2 and 3 and the same wines are served again, in the same order. However, the information given is about the price of the wines rather than their rating. It should be noted that wine C has a clearly better score and is more expensive than wine B. The purpose of this experiment is to examine the extent to which building similar expectations may help reduce gaps in evaluation between two different wines.

The third experiment follows a different framework as it involves two different white wines and one single red wine. The participants first taste two white wines colored in red (rounds 1 and 2). Then the three identical red wines are served one after another, respectively from a decanter – blind (round 3), a bottle of a 40-euro Ghemme (an appellation in the Northern Piedmont) (round 4), and a bottle of a 150-euro Barolo (round 5). The three bottles from which the actual wine came are strictly identical and were opened at the same time. This settings puts the experimenter at risk, given that participants might spot that the exact same wine is served three times. In order to reduce the risk and to try manipulate the preferences of the participants, the names of the first three wines are revealed once they have been tasted (i.e., right after round 3) and the participants are told that the last two wines are meant to conclude the semester in a nice way and that they would therefore not be drunk blind.

3. Empirical analysis

3.1. Effect of expectations on the evaluation (experiment 1)

The results of Experiment 1 are reported in Table 2. Panel A1 shows the ratings of the various wines. More than 80% of the respondents preferred wines 4 and 5 to wines 2 and 1 (which were actually the same). The average scores increase from 70.7 and 76.1 to 83.8 (+13.1) and 86.7 (+10.6) respectively. These differences are significant at the 99% level. Moreover, the median difference in scores is very close to the mean. All this suggests that despite the small sample, the differences are both statistically and economically significant. The prices (panel B1) are totally consistent with the ratings, with one nuance. Giving information on the ratings of wines 1 to 3 reduces the standard deviation of the ratings but increases that of the prices. This can probably be explained by two elements: the information given to the participants pertains to ratings of the wines, and the upward revisions of respondents' expectations are mechanically bounded for ratings (maximum 100 points) but not for prices.

< Insert Table 2 around here >

One potential issue lies in the fact that respondents probably do not all use the same scale to evaluate the ratings and prices of the various wines. Some are more generous, others less. Some maintain small differences between the best and worst wines, while others use a wider range. Finally, some may be extreme when a wine matches their preferences perfectly, or not at all. In short, the distribution of ratings and prices certainly deviates from normal and is potentially very different from one respondent to another. This problem is reinforced by the small sample size. One solution is to use the five-wine ranking. Thus, for each respondent, the original rating (and price) is re-expressed as a ranking from 1 (= highest) to 5 (= lowest). The results reported in panels A2 and B2 are based on the rankings. We can see that the results remain very stable, which again suggests that they are robust.

In order to deepen the analysis, we further examine whether variables associated with individual participants could help explain their reactions. We consider the coherence of participants (i.e., whether their variation in score is coherent with the variation in price), gender, and their wine knowledge self-assessment. As the sample is small, the results of multivariate regressions lack statistical significance. However, a simple t-test shows that "good tasters" tend

to adapt their ratings and prices more strongly after information about the quality and price of wines 1 to 3 has been shared with them. The results are statistically significant for the pair wine 2 vs. wine 4 (differences in ratings: t-stat of -2.17 and p-value of 0.02; differences in prices: t-stat of -1.99 and p-value of 0.04). For the pair wine 1 vs. wine 5, the difference is too small to be significant. This difference between the two pairs may be due to the fact that the information was given just before wine 4 and therefore some participants potentially adapted their expectations more strongly for wine 4 than for wine 5.

3.2. Manipulating expectations to reduce differences in evaluation among different wines (experiment 2)

Figure 1 shows how participants compare wines B (rated 90 points by TWA and sold for about CHF 50) and C (96 points, CHF 70) according to two factors: (1) whether they are in the group that received correct information or not, (2) whether the information concerns the rating of the wine or its price.

< Insert Figure 1 around here >

The left panel of Figure 1 reports the results when information about the ratings of the wines is provided to the participants. Half of them (group 1) taste wine B in round 2 and wine C in round 3, and get correct information. The other half (group 2) tastes the wines in reverse order and thus gets wrong information. The wine served in round 3 obtains a much better ranking both in terms of quality and willingness-to-pay. The result is similar for both groups, which is quite remarkable since the two wines are served in reverse order for group 2 as compared to group 1. When group 1 receives the highest rated and most expensive wine, group 2 receives the lowest rated and least expensive wine, but the information given is the same for everyone: "here is a wine with a TWA rating of 96 points". It must be noted that group 2 reacts less strongly to the information given and that the difference in ranking between the wines served in rounds 2 and 3 is less pronounced for this group, but this group nevertheless shows a preference for the latter wine even though this wine is actually less good (or rather less well rated) and less expensive than the previous one.

The right panel of Figure 1 reports the results when the information provided pertains to the price of the wines tasted. Results are less marked as compared to the ones discussed above.

The group that receives the correct information continues to rate more favorably the wine that is actually the most expensive. The group that receives false information gives a lower rating to the supposedly more expensive wine (but which is actually less expensive) but is still, surprisingly, willing to pay a little more for it. This weaker reaction and the gap between rating and willingness-to-pay for group 2 can certainly be explained by the fact that the prices announced for both wines are in a rather similar range (CHF 50 and CHF 70), whereas their ratings are more dissimilar (90 points, i.e. a very good wine, versus 96 points, i.e. a truly extraordinary wine).

< Insert Table 3 around here >

Table 3 reports the results of a regression the purpose of which is to better understand and assess the statistical significance of the reaction of the participants to the information shared with them. The difference in ratings and willingness-to-pay between wines B and C (dependent variable) is regressed on a set of dummy variables to control if the information shared with the participants is wrong ("WI" in the table) and if it pertains to the price of the wine (IP), and if the participant evaluates himself/herself as a good taster (SA). In order to control for the fact that a good taster may react differently as compared to a taster with less expertise, we also consider a specification in which the wrong information dummy and the self-assessment variables interact together (SA $\times$ WI). In order to ease the interpretation of the results, the table is structured in two panels: Panel A reports the results from the regression, and Panel B analyses their implications.

The intercept shows that participants give generally higher ratings and are ready to pay more for wine C (which is the supposedly better and more expensive wine) than wine B. Respondents who received false information, however, react in a markedly different manner. They tend to prefer wine B to wine C. The corresponding dummy coefficient is statistically significant in all four regressions. Panel B further shows that the difference in ratings and willingness-to-pay between the two wines remains smaller as compared to the participants who received correct information. This means that altering tasters' expectations by giving them false information is sufficient to reverse their ranking (as compared to the group who received true information) but not to completely smooth out the differences between the two wines. Sharing

information about the price of a wine rather than its rating leads to a slightly stronger difference between wines B and C, but the corresponding coefficient is not significant.

The coefficient associated with the taster's self-assessment of their expertise shows that respondents who consider themselves to be good connoisseurs seem to have a stronger preference for wine C as compared to wine B than the other respondents. If we take into account a possible interaction between the self-assessment and the "wrong information" variables, we find that good connoisseurs rated wine C higher than wine B independently of whether the information they received is correct or wrong. It thus seems that they are less sensitive to the release of (wrong) information. This might be due to the fact that they are more confident in their ability to accurately assess a wine. When it comes to their willingness-to-pay, the picture is however a bit different: good connoisseurs who received wrong information are still willing to pay more for wine C but the difference becomes much less substantial. All in one, this discussion suggests that more knowledgeable tasters do not hesitate to express stronger opinions. This observation is consistent with Experiment 1.

3.3. Manipulating expectations to induce differences in evaluation among similar wines (experiment 3)

Figure 2 shows the general results of Experiment 3. We can see that wines C1 (served in round 4) and C2 (round 5) are rated much higher than wine C (round 3), which is identical. The "collective reputation" effect (here Piedmont versus Languedoc) on tasters' expectations seems to be important. We can also observe a difference between wines C1 and C2, but it remains quite small. Both wines come from the same region but the latter benefits from a higher individual reputation. Overall, presenting the same wine in a more prestigious bottle leads to a better evaluation and higher willingness-to-pay. It is interesting to contrast these observations with the fact that, in practice, wine consumers do not necessarily display a preference for more expensive wines. Indeed, Goldstein et al (2008) show that "unless they are experts, [individuals] enjoy more expensive wines slightly less."

< Insert Figure 2 around here >

< Insert Table 4 around here >

Table 4 shows the results of a more detailed analysis of the results of Experiment 3 with a regression. First, it should be noted that the R-squared is low. This is due to the fact that few variables are significant. The results again support the claim that respondents prefer wines C1 and C2 as compared to wine C. The difference in evaluation among the wines is slightly more pronounced (though the difference is not significant) when looking at ratings instead of willingness-to-pay. The coefficient associated with the self-assessment of expertise variable is clearly negative, but significant only in one specification. This suggests that participants who consider themselves good tasters adjust their ratings and willingness-to-pay more strongly. One may further note that participants have a slightly more favorable rating for wine C2 relative to wine C1. The last specification suggests, however, that this difference in ratings is larger for good tasters. In other words, people who rate their knowledge as good find a more substantial difference between wines C1 and C2 than novices. This result is not statistically significant, so we analyze it with caution. Our interpretation is that connoisseurs feel more compelled to revise their evaluation upwards when a visibly more expensive and in principle better wine is served to them. This observation requires further analysis, but it suggests that the manipulation works well and that it works at least as well on more experienced as on less experienced tasters.

4. Conclusions

Our study contributes to the existing literature along several lines. We show that tasters' expectations can be modified and that this often has a significant impact on the rating and willingness-to-pay of tasters. We further show that by using some deception and/or manipulation, one can alter the way participants rate the wines they taste. This leads to changes in ranking, and these changes seem to be more important for tasters with a good level of wine knowledge than for novices. In general, those whose self-evaluation identified them as "good tasters" tend to adapt their ratings and prices more strongly to the information given. That is, one of the differences between novices and more skilled consumers is that the latter express stronger opinions. This can be explained in two ways: (i) they express more assertive opinions because they feel more experienced, and (ii) they may also feel the need to demonstrate that they actually know wine by aligning their opinion with what they think is correct.

Overall, our results suggest that it is possible and relatively easy to manipulate wine consumers into preferring a predetermined wine. It would be good to deepen the analyses presented in this paper. In particular, repeating similar experiments with different and larger samples may allow making the results more robust. In addition, it might be interesting to

examine in more detail whether certain demographic and cultural variables affect the way tasters revise their expectations and evaluation.

5. References

Aqueveque, C. (2018). Ignorant experts and erudite novices: Exploring the Dunning-Kruger effect in wine consumers. *Food Quality and Preference, 65*, 181-184.

Ashton, R. H. (2017). Dimensions of Expertise in Wine Evaluation. *Journal of Wine Economics, 12*(1), 59-83.

Capehart, K. (2019). Does Blind Tasting Work? Another Look. *Journal of Wine Economics, 14*(3), 309-320.

Cardebat, J.-M., & Livat, F. (2016). Wine Experts' Rating: A Matter of Taste? *International Journal of Wine Business Research, 28*(1), 43-58.

Goldstein, R. (2019). Half-Blind Tasting: A Deception-Free Method for Sizing Placebo and Nocebo Responses to Price and Packaging Attributes. *Journal of Wine Economics, 14*(3), 321-331.

Goldstein, R., Almenberg, J., Dreber, A., Emerson, J., Herschkowitsch, A., & Katz, J. (2008). Do More Expensive Wines Taste Better? Evidence from a Large Sample of Blind Tastings. *Journal of Wine Economics, 3*(1), 1-9.

Kruger, J. M., & Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology, 77*, 1121-1134.

Masset, P., Weisskopf, J.-P., & Cossutta, M. (2015). Wine tasters, ratings, and en primeur prices. *Journal of Wine Economics, 10*(1), 75–107.

Niimi, J., Danner, L., & Bastian, S. E. (2019). Wine leads us by our heart not our head: emotions and the wine consumer. *Current Opinion in Food Science, 27*, 23-28.

Niimi, J., Danner, L., Luxing, L., Bossan, H., & Bastian, S. E. (2017). Wine consumers' subjective responses to wine mouthfeel and understanding of wine body. *Food Research International, 99*(1), 115-122.

Parr, W. (2019). Demystifying wine tasting: Cognitive psychology's contribution. *Food Research International, 124*, 230-233.

Parr, W. V., White, K. G., & Heatherbell, D. A. (2003). The nose knows: influence of colour on perception of wine aroma. *Journal of Wine Research, 14*(2-3), 79-101.

Plassmann, H., O'Doherty, J., Rangel, A., & Shiv, B. (2008). Marketing actions can modulate neural representations of experienced pleasantness. *Proceedings of the National Academy of Sciences PNAS, 105*, 1050–1054.

Rodrigues, H., & Parr, W. (2019). Contribution of cross-cultural studies to understanding wine appreciation: A review. *Food Research International, 115*, 251-258.

Siegrist, M., & Cousin, M.-E. (2009). Expectations influence sensory experience in a wine tasting. *Appetite, 52*(3), 762-765.

Spence, C. (2020). Wine psychology: basic & applied. *Cognitive Research: Principles and Implications, 5*(22), 1-18.

Spence, C., & Wang, Q. J. (2019). Wine expertise: perceptual learning in the chemical senses. *Current Opinion in Food Science, 27*, 49-56.

Storchmann, K. (2012). Wine Economics. *Journal of Wine Economics, 7*(1), 1-33.

Veale, R., & Quester, P. (2008). Consumer Sensory Evaluations of Wine Quality: The Respective Influence of Price and Country of Origin. *Journal of Wine Economics, 3*(1), 10-29.

Wang, Q. J., & Spence, C. (2019). Drinking through rosé-coloured glasses: Influence of wine colour on the perception of aroma and flavour in wine experts and novices. *Food Research International, 126*, 108678.

Wang, Q., & Prešern, D. (2018). Does Blind Tasting Work? Investigating the Impact of Training on Blind Tasting Accuracy and Wine Preference. *Journal of Wine Economics, 13*(4), 384-393.

Table 1: Design of the experiments

Panel A: Experiment 1

	Round 1	Round 2	Round 3	After Round 3	Round 4	Round 5
Wine served	A	B	C		B	A
Information given	-	-	-	TWA ratings of A, B and C	-	-
Deception involved	No	No	No		No	No

Wine A: Domaine Courbis, Cornas, Les Eygats 2011 (TWA rating: 93)

Wine B: Domaine de la Charbonnière, Châteauneuf-du-Pape, Les Hautes Brusquières 2010 (TWA rating: 97)

Wine C: Domaine de l'Horizon, Roussillon, Rouge 2008 (TWA rating: 93)

Panel B: Experiment 2

	Round 1	Round 2	Round 3	Round 4	Round 5
Wine served	A	B (group 1) C (group 2)	C (group 1) B (group 2)	B (group 1) C (group 2)	C (group 1) B (group 2)
Information given	-	Rating of B	Rating of C	Price of B	Price of C
Deception involved	No	Yes (*)	Yes (*)	Yes (*)	Yes (*)

(*): group 2 only

Wine A: Clos du Caillou, Côtes du Rhône, La Réserve 2015 (TWA rating: 91, Price: CHF 35)

Wine B: Domaine de Ferrand, Châteauneuf-du-Pape 2015 (TWA: 90, CHF 50)

Wine C: Domaine de la Janasse, Chateauneuf-du-Pape, Chaupin 2015 (TWA: 96, CHF 70)

Panel C: Experiment 3

	Round 1	Round 2	Round 3	After Round 3	Round 4	Round 5
Wine served	A	B	C		C1	C2
Information given	-	-	-	Label of A, B and C	Bottle and label of C1	Bottle and label of C2
Deception involved	Yes (*)	Yes (*)	No		Yes (**)	Yes (**)

(*): white wine coloured in red

(**): the bottle presented does not correspond to the wine that is effectively served

Wine A: Valentina Andrei, Valais, Roussane Marsanne 2021

Wine B: Sybille Kuntz, Mosel, Riesling trocken 2021

Wine C: Mas Jullien, Terrasses du Larzac, Autour de Jonquieres 2016

Wine C1: Mas Jullien, Terrasses du Larzac, Autour de Jonquieres 2016 (served in a bottle of Cantalupo, Ghemme, Collis Breclemae 2011)

Wine C2: Mas Jullien, Terrasses du Larzac, Autour de Jonquieres 2016 (served in a bottle of G. Mascarello, Barolo, Monprivato 2013)

Table 2: Results of experiment 1

Panel A1: original scores

	Wine 1	Wine 2	Wine 3	Wine 4	Wine 5	Delta 4 - 2	Delta 5 - 1
Average	70.71	76.10	62.19	86.74	83.81	10.65***	13.1***
Median	75.00	80.00	75.00	90.00	89.00	10.00	15.00
Std. Deviation	18.09	13.64	27.20	10.74	13.19	11.57	12.72
>0						84%	81%

Panel B1: original log(prices)

	Wine 1	Wine 2	Wine 3	Wine 4	Wine 5	Delta 4 - 2	Delta 5 - 1
Average	3.31	3.67	3.39	4.09	3.99	0.42***	0.68***
Median	3.40	3.81	3.56	4.14	4.01	0.37	0.56
Std. Deviation	0.44	0.42	0.61	0.51	0.51	0.46	0.50
>0						84%	97%

Panel A2: rankings scores

	Wine 1	Wine 2	Wine 3	Wine 4	Wine 5	Delta 4 - 2	Delta 5 - 1
Average	3.82	3.27	4.18	1.55	2.18	-1.73***	-1.65***
Median	4.00	3.00	5.00	1.00	2.00	-2.00	-2.00
Std. Deviation	0.92	0.92	0.92	0.92	0.92	0.92	0.92
<0						84%	81%

Panel B2: rankings log(prices)

	Wine 1	Wine 2	Wine 3	Wine 4	Wine 5	Delta 4 - 2	Delta 5 - 1
Average	4.45	3.06	3.94	1.55	2.00	-1.52***	-2.45***
Median	4.50	3.00	4.00	1.00	2.00	-1.50	-3.00
Std. Deviation	0.61	1.05	1.05	0.72	0.85	1.41	1.06
<0						84%	97%

*Note: *, ** and *** denote significance at the 90%, 95% and 99% level respectively.*

Table 3: Regression analysis of experiment 2

Panel A: Regression results

	Wine C vs. Wine B: Difference in ratings		Wine C vs. Wine B: Difference in willingness to pay	
	(I)	**(II)**	**(III)**	**(IV)**
Intercept	-0.82	-1.24	-1.25***	-1.54***
	(0.66)	(0.79)	(0.44)	(0.53)
Wrong information (WI)	1.31**	1.87**	2.43***	2.8***
	(0.64)	(0.86)	(0.44)	(0.58)
Information on price (IP)	-0.39	-0.35	-0.31	-0.25
	(0.61)	(0.61)	(0.42)	(0.42)
Self-assessment (SA)	-1.64**	-0.94	-1.1**	-0.6
	(0.64)	(0.97)	(0.44)	(0.67)
Interaction SA × WI		-1.27		-0.86
		(1.3)		(0.89)
Nobs.	35	35	34	34
R2	0.33	0.35	0.61	0.62

*Note: standard errors are in brackets. *, ** and *** denote significance at the 90%, 95% and 99% level respectively.*

Panel B: Implications (based on regression coefficients)

|	Wine C vs. Wine B: Difference in ratings		Wine C vs. Wine B: Difference in willingness to pay	
	(I)	**(II)**	**(III)**	**(IV)**
Reaction (difference between wine C and wine B) of all tasters:				
True information on rating	-0.82	-1.24	-1.25	-1.54
True information on price	-1.21	-1.59	-1.56	-1.79
Wrong information on rating	0.49	0.63	1.18	1.26
Wrong information on price	0.10	0.28	0.87	1.01
Reaction (difference between wine C and wine B) of "good" tasters:				
True information on rating	-2.47	-2.18	-2.35	-2.14
True information on price	-2.86	-2.53	-2.66	-2.39
Wrong information on rating	-1.15	-1.58	0.08	-0.21
Wrong information on price	-1.55	-1.92	-0.23	-0.46

Note: This panel shows the implications of the results (coefficients) from Panel A in terms of reaction (i.e., change in rating or price) of the tasters following the release of information about the rating or the price of a wine.

Table 4: Regression analysis of experiment 3

	Wine C1 vs. C	Wine C2 vs. C	Wines C1 & C2 vs. C	
Intercept	-1.4***	-1.62***	-1.36***	-1.41***
	(0.31)	(0.47)	(0.35)	(0.36)
Delta price	0.33	0.37	0.35	0.35
	(0.43)	(0.65)	(0.39)	(0.39)
Self-assessment (SA)	-0.76	-1.31	-1.04*	-0.76
	(0.57)	(0.88)	(0.52)	(0.74)
Wine C2 vs C1 (dC)			-0.29	-0.2
			(0.39)	(0.43)
Interaction SA x dC				-0.55
				(1.05)
Nobs.	23	25	48	48
R2	0.10	0.10	0.11	0.11

*Note: standard errors are in brackets. *, ** and *** denote significance at the 90%, 95% and 99% level respectively.*

Figure 1: Summary results of experiment 2
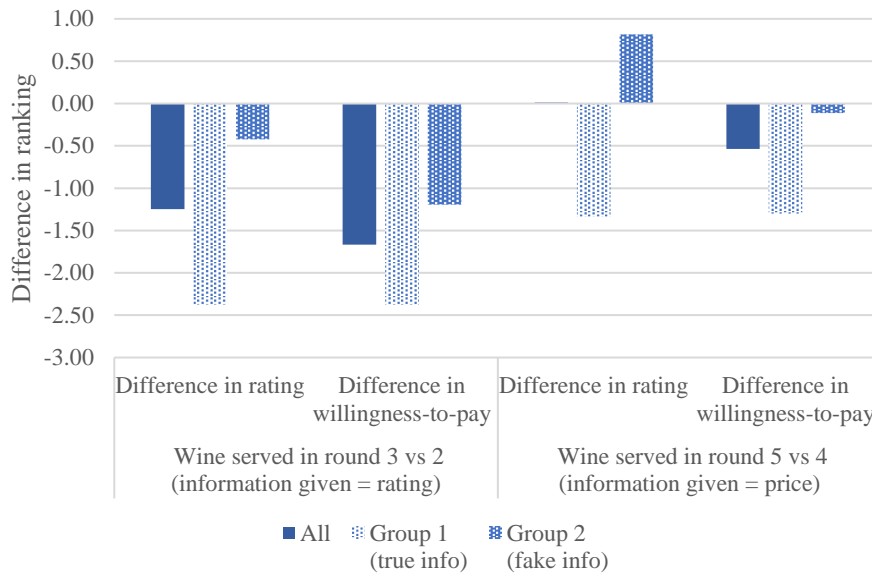
Figure 2: Summary results of experiment 3