Expert Opinion and the Demand for Wine

James Hilger, Greg Rafert, and Sofia Villas-Boas

We examine the impact of expert opinion on retail wine sales utilizing an experiment implemented on a national retail chain. A test store was selected to receive an expert-ranking label for selected wines. A comparable control store was also chosen. Sales effects of the ranking were examined and compared between test and control stores.

Product awareness and perceptions of product quality can have large effects on consumption patterns. Given the variety of methods employed by manufacturers and marketers to inform consumers of a product's quality, recent research has analyzed the extent to which product-quality information affects consumer behavior. This literature examines the effect of a variety of information types and sources, including branding, mandatory product labeling, and advertising. One additional method used to convey quality information to consumers is through so-called experts. For example, *Consumer Reports* tests a large number of products each year and publishes product reviews, and magazines such as *Wine Spectator* and *Wine Enthusiast* rate wine quality.

Most studies analyze the impact of expert opinion on consumer demand for goods for which the quality is only learned by consumers after consumption (the so called "experience goods"). However, these studies face a significant obstacle: products of high quality are likely to both receive high-quality ratings from experts and to be of high quality. As such, it is difficult to determine the extent to which consumer demand is affected by expert reviews, since to do so, the researcher must control for unobservable product quality. Yet even if expert reviews affect consumer demand for a particular good, demand may change because consumers respond to the quality signal in the review or alternatively, because consumers are merely alerted to the presence of that good.

Experimental Design and Data

The research goal is to examine the impact of expert opinion on retail wine purchases. To distinguish the effect of expert reviews from that of product quality, we utilize an experimental approach implemented at stores in a national retail grocery chain that cooperated with the research team. Wines in a retail store in Northern California were randomly chosen to display wine scores from a proprietary wine scoring system, and wine opinion labels were then displayed for one month during the spring of 2006. The retailer classifies the chosen treatment store as a high wine-revenue store, and the store has wine revenues that are greater than the revenues for most other stores operated by the retailer in California. Further, on average, the store is located in a wealthier area, has a greater amount of shelf space dedicated to the sale of wine, stocks more expensive wines, and sells more wine as a percentage of total grocery sales. To the extent that consumers in more wealthy areas and those buying more expensive wines are likely to be more fully informed regarding wine quality than consumers in other areas, we have selected a store that should reduce the likelihood of finding a significant treatment effect.

Wine scores from a proprietary wine scoring system were displayed in the treatment store for four weeks during the month of April 2006 for a random selection of wines. The wines chosen for the experiment were not selected from the total population of wines in the store since many wines do not receive wine scores from any of the wine-rating agencies. Instead, the wines were chosen from the population of wines stocked in the store that received wine scores. Of the total of 1,089 wines sold in the test store in March 2006, 476, or 44 percent, received wine scores from one of several potential wine-scoring agencies. Thus, by selecting 150 treatment wines, we treated 32 percent of the total population of potential candidates and 14 percent of all wines within the store.

To each treated wine, we affixed a label to the shelf below that indicated the score awarded the wine from the scoring system. Each label displayed information on the score received by a wine, the wine's price, as well as the name of the proprietary scoring system. Wine scores awarded by the scoring system can in theory range from 50 to 100, with 100 being the highest possible score. In practice, however, wine scores typically range between 75 and 100, with most wines receiving scores between 80 and 89.

We obtained weekly store-level sales data from the grocery chain for each wine sold in all Northern Californian stores. The data provided information on the number of bottles sold, the pre-discount price, the discount amount, and the wine variety. The weekly sales data were aggregated to the month-level for each store to generate total number of bottles sold per month, average pre-discount price, average post-discount price, and whether a bottle of wine was discounted in any one week during a given

month. For those wines for which wine scores exist, we then merged wine score information from the proprietary wine score system with the wine sales data.

Due to differences between the retail chain's database of stocked wines and those wines actually stocked at the time of the experiment within the retail store, 112 wines were labeled in the test store. There are few differences between treated wines and untreated wines for which scores exist, as can be seen in Table 1. For example, the mean score for treated wines is equal to 84.1 while the mean score for untreated wines with scores is 83.7, and this difference is not significant. Further, the pre-treatment difference between price and quantity is not significantly different for these groups, thereby suggesting that the selection of the treatment wines was random. There are also not significant observable differences between treated wines and untreated wines for which scores are not available.

Empirical Strategy and Findings

Given the experimental design, we utilize a differences-in-differences approach to analyze the effect of the treatment on treated wines and to determine whether expert opinion provided quality information or simply highlighted the existence of treated wines. Specifically, we first examine the effect of the treatment on the treated wines by comparing the change in the sales of treated wines from the pre-treatment to treatment month in the test store, to the change in the sales of treated wines from the pre-treatment to treatment month in the control store.

Figure 1 illustrates the basic idea behind the difference-in-difference approach to identify the impact of the treatment via expert opinion labeling. Consider on the left, two bars corresponding to the number of bottles sold in store C (that was the store that did not change the way the products were displayed on the shelves), where the first bar corresponds to the before period, and the second bar corresponds to the bottles sold in C in the after period. In the middle of Figure 1, let us represent store T (the test store where we displayed expert opinion labels). The two bars correspond to the number of bottles sold in the before and in the after period, respectively. For store C, we perform a first difference consisting of the changes in bottles sold from the after period relative to the before period, and call that D_C. Next we do the same difference of after minus before sales for the test store, and call that difference D_T. The effect of our treatment consists then of the difference D_T-D_C, that is, the difference in these two differences. And in doing so, we assess the changes in the test store relative to the changes in the control store.

We run the above analysis first on only those wines that received an expert opinion label. The dependent variable is the number of bottles of wine sold of a product in a store in a certain week, and the independent (explanatory) variables are (i) an indicator variable, *store,* that is equal to one for treated wines in the test store and equal to zero for treated wines in the control store, (ii) an indicator variable, *month*, that is equal to one during the treatment month and equal to zero during the pre-treatment month. The coefficient on *store* can be interpreted as a treatment group-specific effect, the one on *month* as a time trend common to the control and test stores, and the *(store*month)* coefficient can be interpreted as the true effect of the treatment. We control for potentially important other factors (covariates) such as promotions or discounts which, if omitted, could lead to a biased estimate of the treatment effect.

The average effect of the treatment on the treated wines is not significantly different from zero. The only variable which is significant is the promotion variable. It is always positive, indicating that a wine placed on promotion sometime during the month (where the minimum promotion length in the data is two weeks and the maximum is four weeks during a month) can expect on average to sell approximately 13 to 15 bottles more per month than if it were not discounted. Since non-promoted, treated wines sold an average of four bottles, this effect indicates that the average number of bottles sold of a treated wine increases by 425 to 475 percent when it is placed on promotion.

Although useful for examining the average treatment effect on the treated, the above investigation does not address the extent to which the expert opinion effect is related to quality information provision versus general publicity. To examine the manner in which consumers use expert opinion information, we include interactions between score, price, and the treatment. If expert opinion primarily provides quality information to consumers, then only those treated wines that received higher scores should experience an increase in quantity sold. Alternatively, if the primary effect of

Table 1. Descriptive Statistics

	Treated Wines	Untreated Wines (With Scores)	Untreated Wines (Without Scores)
Score	84.1	83.7	
	[3.5]	[3.0]	
Quantity March (pre)	12.2	14.3	9.2
	[20.3]	[19.9]	[18.2]
Quantity April (post)	14.5	18.4	9.1
	[21.9]	[20.0]	[18.0]
Price (pre)	11.8	10.9	11.8
	[7.8]	[6.3]	[9.0]
Price (post)	12.5	11.6	11.9
	[10.3]	[7.2]	[8.9]
Percent Red	63.4	61.9	60.6
Number of Observations	**112**	**253**	**629**

Source: Retailer provided scanner data set.

expert opinion labels is to alert consumers to the existence of a wine, then the treatment should have an impact irrespective of a wine's score. As above, to reduce the likelihood that the estimated treatment effects are biased, we include price and whether a wine was discounted in any one week during a given month, as well as interactions between price, score, discount, and the treatment.

We find that although there is no overall consumer response to expert opinion provision, a subset of highly reviewed wines experienced an increase in demand. In particular, the estimated effect on a treated, low-price wine of moving from low to high score lies between eight and 15 additional bottles sold during a given month. Given that low-priced, high-scoring wines sold an average of 26 bottles during March in the test store, sales increased by an average of 30 to 58 percent as a result of the treatment.

Although we are primarily interested in estimating the average treatment effect on the treated wines, we also estimate the average treatment effect on the untreated wines. Interestingly, we also find that as demand increased for a subset of treated wines, demand did not change for untreated wines. Thus, consumers either did not completely substitute towards treated wines or a sufficient number of consumers entered into the wine market to offset those consumers who substituted away from untreated wines.

Conclusions

Our results strongly suggest that expert opinion can affect the demand for wine by transmitting product-quality information to consumers. Results indicate that consumers utilize quality information provided by expert opinion labels, as opposed to solely using the label to learn of a wine's existence. Unlike most previous work that examines the impact of expert opinion on consumer demand, we are able to disentangle the endogenous relationship between product quality and expert opinion provision through the use of an experimental approach in a large national retail grocery chain. We randomly select 150 wines to display expert opinion information. Then we select a control store with similar characteristics to those of the test store. We are then able to examine both the effect of expert opinion on the overall demand for wine, and the role of expert opinion labels in providing quality information versus alerting consumers to the existence of a wine.

We find that on average, sales of wines with expert opinion information did not increase. However, we do show that low-priced, high-scoring wines experienced an increase in demand relative to other treated wines. These results are robust to the use of alternate control stores, the use of the alternate test store, and the variables included within the regressions. Further, these effects only exist during the treatment period, and are not found when other pre-treatment months are used as the treatment period. Although we can offer no definitive evidence, one potential explanation for the lack of a high-score effect for more expensive wines is that consumers who purchase expensive wines are more fully informed regarding product quality, and thus gain little information when expert opinion is displayed. Finally, we find that as demand increased for a subset of treated wines, demand did not change for untreated wines.

Figure 1. Difference in Difference

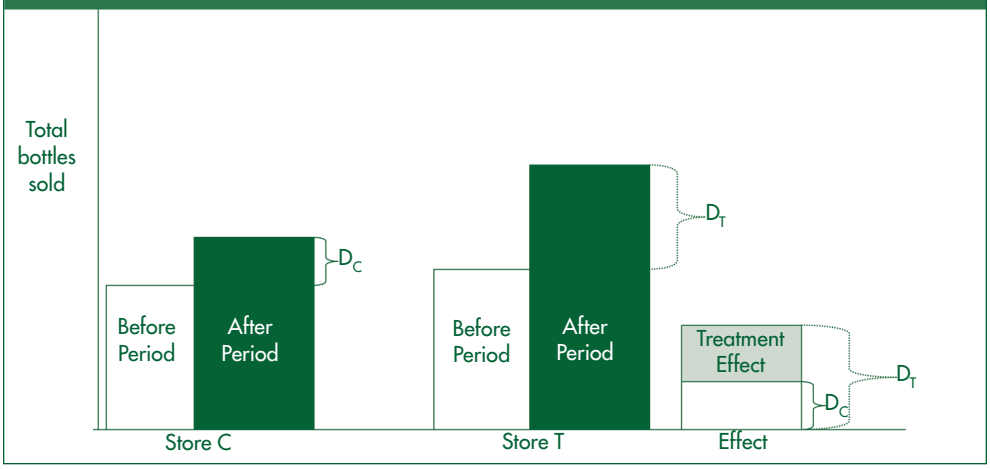

Our findings broadly suggest that expert opinion can provide quality information to consumers and that at least some consumers will use such information when making purchasing decisions. To the extent that certain consumers previously did not participate in the market due to a lack of product information, such information provision may allow the market to expand as new consumers enter. Further, as quality information is distributed and consumers learn which producers are associated with high-quality products, low-quality producers may increase their product quality to more effectively compete with high-quality producers. Both the relationship between information provision and consumer entry, and that between quality information and the quality provided by producers remain as interesting avenues for further research.

James Hilger is an economist in the Bureau of Economics, Federal Trade Commission, who can reached at hilger@ftc.gov. Greg Rafert has a Ph.D. from ARE and Sofia Villas-Boas is an associate professor, both in the ARE department at UC Berkeley. They can be reached by e-mail at rafert@are.berkeley.edu and sberto@are.berkeley.edu, respectfully. *The authors gratefully acknowledge the financial support received from the Giannini Foundation.*