

Expert Opinion and the Demand for Experience Goods: An Experimental Approach in the Retail Wine Market

James Hilger, Greg Rafert, Sofia Villas-Boas*

November 16, 2007

Abstract

The effect of expert opinion on consumer demand for experience goods is difficult to quantify as the relationship between reviews and purchases may be driven by product quality, and further, it is unclear whether expert opinion increases demand through the provision of quality or existence information. Utilizing an experimental approach in the retail wine market to overcome these obstacles, we find that although there is no overall consumer response to expert opinion labels for wine, demand for a subset of highly reviewed wines increased; indicating that labels transmit quality information, as opposed to solely informing consumers of a wine's existence.

*Hilger: Bureau of Economics, Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington DC 20580; Rafert and Villas-Boas: Department of Agricultural and Resource Economics, University of California, Berkeley, CA, 94720. Hilger is an economist at the Federal Trade Commission. The views expressed in this paper are those of the author and do not necessarily represent the views of the FTC, any individual Commissioner, or any Bureau. We thank Stefano DellaVigna and Daniel Hosken for helpful comments. We are also thankful to Michelle Scizak, Reed Johnson, and Steve Flaming for help with the data; Will Liu and Jeff Coronado for help with the experimental implementation; and Kyle Birchard, Elizabeta Perova, Grant Chen, Patricia Javier, Katan Patel, and Elizabeth Creed for their excellent research assistance. The authors gratefully acknowledge the financial support received from the Giannini Foundation. Corresponding author: James Hilger at jhilger@ftc.gov.

1 Introduction

Product awareness and perceptions of product quality can have large effects on consumption patterns. As a result, manufacturers and marketers have developed a number of methods to both increase product awareness and to broadcast product quality to potential consumers. The methods employed to inform consumers about product quality are particularly important for experience goods, since consumers may only fully determine quality after purchase.¹

Given the pervasiveness of experience goods within the marketplace, there exists a large and growing theoretical literature that examines ways in which uncertainty regarding product quality affects consumer demand (see Akerlof, 1970; Nelson, 1970; Wiggins and Lane, 1983; and Wolinsky, 1995). Further, given the variety of methods employed by manufacturers and marketers to inform consumers of a product's quality, a closely related empirical literature has developed that analyzes 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 (Montgomery and Wernerfelt, 1994), mandatory product labeling (Jin and Leslie, 2003), and advertising (Akerberg, 2001; Akerberg, 2003).

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, the Zagat guide gives quality ratings to restaurants in U.S. metropolitan areas, Ebert and Roper review movies, and magazines such as Wine Spectator and Wine Enthusiast rate wine quality.

Most studies that analyze the impact of expert opinion on consumer demand for experience goods 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. To our knowledge, there exist only two studies that attempt to isolate the impact of expert reviews and product quality on consumer demand. Eliashberg and Shugan

¹Examples of experience goods include automobiles, restaurants, movies, and wine.

(1997), who examine the effect of movie critics on the demand for movies, find that movie critics appear to have little effect on consumer demand. Reinstein and Snyder (2005) also focus on the motion picture industry by exploiting the timing of movie reviews by Siskel and Ebert. The authors find no overall effect of reviews, but show that positive reviews increased box office revenues for narrowly-released movies and dramas. Although their identification strategy more convincingly isolates the effect of expert opinion from product quality than that used in Eliashberg and Shugan (1997), it is unclear why demand increased for only narrowly-released movies and dramas, and not other films.

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. We are aware of only two papers that investigate the extent to which any publicity is good publicity. In their analysis of the impact of positive and negative book reviews in the New York Times, Sorensen and Rasmussen (2004) show that although both positive and negative reviews increase book sales, positive reviews have a larger effect on book sales than negative reviews. Reinstein and Snyder (2005) find similar results that indicate that only positive movie reviews affect movie demand.

Our paper contributes to the empirical literature by examining 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. 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. Based on wine sales trends for previous years, a control store was subsequently selected for the treatment store to allow for the use of a difference-in-difference approach. We then tested whether consumers responded to expert opinion, and investigated the extent to which any publicity is good publicity by examining consumers' responses across wines of differing quality.

We find that on average, sales of wines with expert opinion information did not increase. However, we do find that demand for a subset of the treated wines

increased. In particular, low-priced high-scoring wines experienced an increase in demand relative to other treated wines. Consumers therefore only responded to high quality information signals and only for a subset of wines, suggesting that not all publicity is good publicity. To the extent that consumers who purchase expensive wines are more fully informed regarding product quality, the effect of expert opinion provision should diminish. Thus, these results may indicate that the demand for higher scoring expensive wines did not increase because consumers were already sufficiently informed. Alternatively, our results may be indicative of the length of time required for consumers to fully adopt and trust a new source of information. Instead of fully trusting the new information source immediately, consumers may attempt to verify the accuracy of the new source by sampling labeled wines for which the costs of experimentation are low. That is, they may only purchase wines that are relatively inexpensive. To the extent that this explanation is valid, it suggests that a treatment period of one month may not be sufficient to observe the full effect of expert opinion provision.

Interestingly, we also find that as demand increased for a subset of treated wines, demand did not change for untreated wines. There are several potential explanations for this last finding. First, it may be that substitution by consumers towards treated wines and away from untreated wines was not one-for-one. That is, at least some consumers, when buying labeled wines, also continued to buy unlabeled wines. Alternatively, substitution may have been one-for-one, but consumers who previously did not purchase wine due to a lack of quality information entered into the wine market when expert opinion information was provided. Finally, consumers may have substituted temporarily by stocking up on treated wines or spatially by reducing the quantity of wine purchased at competing stores.

The remainder of the paper is structured as follows. Section 2 gives a simple theoretical model. Section 3 discusses the design of the experiment, and Section 4 gives our empirical strategy. Section 5 presents results, and Section 6 concludes.

2 Theoretical Framework

We illustrate the potential impacts of expert opinion in the wine market with a simple model of consumer demand. The model shows that expert opinion provision does not necessarily increase demand for wine. Further, it indicates that both positive and negative wine reviews may increase wine demand for a given consumer.

Let k be an individual's familiarity with the existence of a given bottle of wine and l be that individual's perception of the wine's quality. We assume that an individual buys a bottle of wine if two conditions are satisfied: (1) $k \geq k_{min}$ and (2) $l \geq l_{min}$. That is, an individual must have a minimal level of knowledge regarding the existence of the wine (condition 1) and the consumer's perception of the wine's quality level must exceed a minimum threshold (condition 2).

Both positive and negative expert reviews for wine increase k . For simplicity, we assume that if a wine is not reviewed then k falls below k_{min} , and if a wine is reviewed, k necessarily exceeds k_{min} .² We make this assumption because in a retail store that stocks a large number of wines, a consumer is unlikely to notice a given bottle of wine, unless the bottle is highlighted by, for example, an expert opinion label. One effect of a review is therefore to highlight the existence of a wine and to increase the likelihood that a consumer purchases that wine. Although a wine review is a necessary condition for a wine to be purchased in this simple model, it is not sufficient since l , an individual's perception of a wine's quality, must exceed l_{min} . Denote l_{PR} as the value l takes when a wine receives a positive review and l_{NR} as the value l takes when a wine receives a negative review. We analyze three cases below.³

If $l_{PR}, l_{NR} < l_{min}$, then expert reviews do not affect wine demand. This may occur if a consumer remains skeptical of a wine's quality even after a positive review. For example, a wine from an unknown wine growing region, or a wine made from a varietal for which the consumer has little past experience may deter a consumer

²The model's results are not qualitatively affected by relaxing this assumption.

³To simplify the discussion we do not consider the case where a poor review may decrease consumer demand. Although this is likely, including such a possibility does not alter the conclusions derived from our model.

from purchasing that wine. When $l_{PR}, l_{NR} \geq l_{min}$, positive reviews have the same effect on the probability of purchase as negative reviews. To the extent that the reviewer lacks credibility with the consumer or the negative review merely allows the consumer’s perception of a wine’s quality to exceed l_{min} , the information content of negative reviews may be largely ignored by consumers. In this case, any publicity is good publicity and reviews serve to highlight a wine’s existence. Reviews only have a differential impact on wine demand when $l_{PR} \geq l_{min}$ and $l_{NR} < l_{min}$. That is, when a positive review allows a wine to exceed the minimum necessary quality level, and when a negative review would not allow a wine to exceed the minimum quality level, only positive reviews will have an impact on wine demand. In this case, only positive reviews are influential in determining wine purchases. In what follows, we investigate both the highlighting and informative effect of expert opinion.

3 Experimental Design

To distinguish the effect of expert reviews from that of product quality we utilize an experimental approach at a large national retail grocery chain in which we randomly select 150 wines stocked in two retail stores.⁴ Although we have two potential treatment stores that can be used in the analysis, the discussion that follows primarily uses only one of the two stores to simplify the analysis. Since both stores have similar characteristics and as we show in Section 5, the reported results are invariant to the treated store utilized in the analysis, the use of only one of the treatment stores does not qualitatively affect our results.

The retailer classifies the chosen treatment store as a high wine revenue store, and as Table 1 indicates, 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

⁴The grocery chain has a large presence in Northern California and obtains substantial revenues from the sale of wine, beer, and liquor. As with many other grocery stores, consumers can choose from a large number of wines located on four to five levels of shelves along both sides of an aisle.

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 treatment 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 the wine that indicated the score awarded the wine from the scoring system. Figure 1 shows a label used in the experiment. As it demonstrates, each label displays information on the score received by a wine, the wine's price, as well as the name of the proprietary scoring system.⁵ 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 provide information on the number of bottles sold, the pre-discount price, the discount amount, and the wine varietal. The weekly sales data are aggregated to the month-level for each store to

⁵Examples of wine rating agencies include the Wine Spectator, Wine Enthusiast, and Wine Advocate. Wineries generally send samples of their wines to each of these agencies where the wines are tasted blind. Wines are either judged by a single individual, or several judges sample each wine and the scores are averaged. In general, if the final score is greater than 75, the score is reported on the agency's website, which requires a paid subscription, and in the agency's monthly magazine or newsletter.

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 merge 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 treatment store. Descriptive statistics for treated wines, untreated wines with scores, and untreated wines without scores are given in Table 2 for the pre-treatment month (March) and treatment month (April) in the treatment store. As the table indicates, there are few differences between treated wines and untreated wines for which scores exist. 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. 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. Table 2 also shows that there are few observable differences between treated wines and untreated wines for which scores are not available since the observable differences between these groups are small.

To rigorously examine the extent to which the treatment increased wine sales, we take advantage of the large number of stores in the dataset to select a suitable control for the treatment store. To reduce the number of potential control store candidates, we first restrict the analysis to other high revenue wine stores. This restriction greatly reduces the number of potential control stores to 38. Given that there exists a high degree of correlation between wine sales, wine selection, and the demographics of areas surrounding such stores, this restriction ensures that the treatment stores are likely to be matched to stores with similar characteristics.

To select one control for the treatment store from the 38 candidates, we use a methodology that aims to ensure that the (1) effect of price, discounts, and wine type on sales of wine are similar across the treatment and control store, and (2) pre-treatment time trends in total number of bottles sold during each month are similar for both the treatment and control store. The latter condition is similar to a robust-

ness check used in many difference-in-difference approaches to determine whether there exist differences in the pre-treatment trends (see for example Meyer, 1995), while the former condition helps to ensure that differential responses to changes in price and the existence of promotions across treatment and control stores are small or nonexistent, thus making it less likely that the estimated treatment effects are biased.

Specifically, for the treatment store we estimate the following equation for the 18 months preceding the treatment intervention:

$$\begin{aligned}
 Q_{it} = & \alpha + \beta_1(\text{price})_{it} + \beta_2(\text{discount})_{it} + \beta_3(\text{red})_{it} + \beta_4(\text{price} * \text{red})_{it} \\
 & + \beta_5(\text{price} * \text{discount})_{it} + \beta_6(\text{red} * \text{discount})_{it} + \beta_7(\text{month})_t \\
 & + \beta_8(\text{month} * \text{price})_{it} + \beta_9(\text{month} * \text{discount})_{it} + \beta_{10}(\text{month} * \text{red})_{it} + \epsilon_{it}
 \end{aligned}$$

where Q_{it} is the number of bottles sold of wine i during month t , $price$ is the average price for wine i during month t , $discount$ is a dummy variable equal to one if a wine was on sale for any one week during month t , red indicates if a given wine is a red wine, and $month$ is a vector of month fixed effects. We use the estimates from the regression to generate predicted quantity values for each candidate control store and we then calculate the difference between the actual and predicted quantity value for each observation in the control store. The prediction residuals within a control store are then utilized to obtain the average and standard deviation of the residuals.

A control store that is similar to the treatment store will have residuals with mean equal to zero. For example, if price, discounts, wine type, and time have similar effects in both the treatment and control store, when the estimates from the above regression are used to generate the difference between predicted and actual quantity values, the residuals should on average be equal to zero. We therefore select the control for the treatment store that has a residual mean equal to 0.3 and a standard deviation of 24.1. This value lies very close to zero and given the standard deviation, we cannot reject the null hypothesis that the mean differs from zero.⁶

⁶The residual mean for all control stores ranges from 0.02 to 12.0 and the standard deviation ranges from 11.1 to 39.9. Although there exist two stores with lower residual means, we use the control store with a residual mean equal to 0.3 because it performs substantially better than the

The selected control store is an appropriate control in that it is similar to the treatment store in two additional respects. First, as Table 1 shows, the store characteristics for the control store differ only marginally from those of the treatment store. Second, the use of an alternate method for selecting the control store yields results that are similar to those obtained using the residual methodology. Specifically, we regress log quantity in the treatment store against log quantity for each potential control store for the 18 months preceding the treatment intervention:

$$\log(Q_{it}^t) = \alpha + \beta_1 \log(Q_{it}^c) + \epsilon_{it}$$

where $\log(Q_{it}^t)$ is the log of quantity for wine i in time t in the treatment store and $\log(Q_{it}^c)$ is the log of quantity in control store c .⁷ In this method, a control store that has an identical pre-treatment trend in total number of wine bottles sold each month will have an estimated coefficient for β_1 equal to one, thereby indicating that a 1 percent increase in quantity in the control store is correlated with a 1 percent increase in quantity in the treatment store. The control store identified by using the residual method described earlier has an estimate for β_1 that is equal to 0.97.⁸ Thus, monthly wine sales in the control store move similarly to those in the treatment store. Given that the trends in monthly sales are very similar, the control and treatment store have similar observable characteristics as identified in Table 1, and the residual methodology indicates that little difference exists between the predicted and actual quantity values, the control store is very likely to be an appropriate control for the treatment store.

other two stores under additional selection criteria described below.

⁷This method closely mimics the robustness test used in many difference-in-difference approaches to determine whether there exist differences in the pre-treatment trends (see Meyer, 1995).

⁸The point estimate for β_1 for all control stores ranges from 0.09 to 0.99. Only one store has a point estimate for β_1 that exceeds that of our selected control store. We use the control store with β_1 equal to 0.97 since this store performs better under the two selection criteria previously discussed.

4 Empirical Strategy

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 treatment store, to the change in the sales of treated wines from the pre-treatment to treatment month in the control store. We do so by running the following difference-in-difference specification for the pre-treatment and treatment month on only those wines that received an expert opinion label:

$$(1) \quad Q_{ist} = \beta_0 + \beta_1 T_{is} + \beta_2 t_{it} + \beta_3 T_{is} * t_{it} + \epsilon_{ist}$$

where Q_{ist} is the number of bottles of wine i sold in store s in time t , T_{is} is an indicator variable that is equal to one for treated wines in the treatment store and equal to zero for treated wines in the control store, and t_{it} is a month dummy that is equal to one during the treatment month and equal to zero during the pre-treatment month. The coefficient on T_{is} can be interpreted as a treatment group specific effect, that on t_{it} as a time trend common to the control and treatment stores, and the coefficient for $T_{is} * t_{it}$ can be interpreted as the true effect of the treatment. This specification corresponds to Specification 1 in the tables below.

Although useful for examining the average treatment effect on the treated, this specification 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 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.

This specification also fails to control for potentially important covariates

such as promotions or discounts, which if omitted, could lead to a biased estimate of the treatment effect. For example, if wines in the treated store were all placed on sale during the treatment month, and these wines were not put on sale in the control store, we would falsely attribute an increase in the sale of such wines to the treatment when in fact much of the increase in the number of bottles sold may have been the result of lower prices. 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, discount, and the treatment.

Finally, the simple difference-in-difference approach does not take into account that there exist many different types of wine and that consumer response to expert opinion provision may differ across wine varieties. For example, those consumers that purchase red wines may on average be more knowledgeable regarding wine quality than those that purchase white wines. To account for this possibility, we include a single dummy variable that indicates if a given bottle is a red wine and interactions between this variable and the treatment.

As a result of these additions, we estimate the following difference-in-difference specification:

$$(2) \quad Q_{ist} = \beta_0 + \beta_1 T_{is} + \beta_2 t_{it} + \beta_3 T_{is} * t_{it} + \delta X_{ist} + \gamma(X_{ist} * T_{is}) + \lambda(X_{ist} * t_{it}) \\ + \theta(X_{ist} * T_{is} * t_{it}) + \pi(\text{price}_{ist} * \text{score}_{ist} * T_{is} * t_{it}) + \epsilon_{ist}$$

where X_{ist} is a matrix that contains the variables *price*, *discount*, *red*, and *score*.^{9,10}

⁹We also estimated models in which the red and discount dummy variables were interacted with T_{it} , t_{it} , *price*, and *score*. Our results do not qualitatively change, and thus we exclude these specifications from our discussion of the results.

¹⁰Although the retailer provides weekly data, we aggregate the data to the month level because in general, the retailer only changes wine prices once a month at the beginning of each month. For the selected treatment and control store, 13 percent and 11 percent of wines, respectively, experience a price change within a given month. Thus, in general, prices for a wine remain constant during a month. Further, of wines that experience a price change, 94 percent change prices in a similar manner in both the treatment and control store. In addition, the only other time varying variable we include, whether a wine is placed on promotion, also generally only changes at the beginning of a given month. Only 8 percent of wines in the treatment store and 2 percent of wines in the control store experience a change in promotional status during the beginning of the second, third,

Price is coded as a dummy variable that is equal to one if the price of wine is less than \$12, *discount* is a dummy variable equal to one if a wine was on sale for any one week during a month, *red* indicates if a given wine is a red wine, and *score* is a dummy variable equal to one if a wine received a score from the proprietary wine rating system greater than or equal to 86. The cutoff point of 86 is chosen since scores lower than 86 are seen by many wine producers as poor scores that are unlikely to increase wine demand, while the cutoff point of \$12 for price is chosen because the price level of \$12 is often used by both wine retailers and in consumer surveys as distinguishing expensive from less expensive wines.^{11,12} This specification corresponds to Specification 2 in the tables that follow.

The primary coefficients of interest are those coefficients in the vector θ and the estimated coefficient for π . The parameters in θ allow us to examine to what extent wine characteristics such as score and price interact with the treatment. For example, if the coefficient of the overall treatment effect is not significantly different from zero and the estimate of the interaction between score, T_{is} , and t_{it} is positive and significant, then we can conclude that the treatment only increased the sales of high scoring wines. Such a finding would support the hypothesis that expert opinion provides quality information and does not simply serve an attention-grabbing role. However, the effect of score may differ across treated wines. If consumers who purchase expensive wines are sufficiently informed regarding wine quality, while consumers who purchase less expensive wines lack knowledge of wine quality, then quality information provision should only affect the demand for less expensive wines. The parameter π allows for such a possibility by permitting the impact of score to vary across wine price.

Although we are primarily interested in estimating the average treatment effect on the treated, we also estimate the average treatment effect on the untreated.

or fourth week of a month. As a result, although our data is provided for each week, there exists very little inter-month variation. Our results are not qualitatively affected by using monthly data.

¹¹The results are not qualitatively affected by the selection of different cutoffs for price or score. Specifically, although the magnitude of the estimates change, the direction and significance of the estimated effects are not affected.

¹²Approximately 66 percent of wines have a price that is less than \$12 and approximately 33 percent of wines have a score that is greater than or equal to 86.

A priori, it is not clear whether this estimate should be less than, greater than, or equal to zero. For example, as consumers purchase wines with expert opinion labels, they may substitute away from unlabeled wines. From a retail grocer’s perspective, this is likely to be quite important since the extent of substitution will affect wine revenues. Alternatively, consumers who previously did not purchase wine due to a lack of information may be induced into entering the market since more information is now available. To determine the extent to which the treatment affected sales of untreated wines we estimate the following difference-in-difference specification on only those wines that did not receive an expert opinion label in the treatment store:

$$(3) \quad Q_{ist} = \beta_0 + \beta_1 T_{is} + \beta_2 t_{it} + \beta_3 T_{is} * t_{it} + \delta X_{ist} + \gamma(X_{ist} * T_{is}) + \lambda(X_{ist} * t_{it}) + \theta(X_{ist} * T_{is} * t_{it}) + \epsilon_{ist}$$

where the matrix X_{ist} now contains only the variables price, discount, and red. This specification corresponds to Specification 3 in Table 10.

5 Results

The columns of Tables 3a and 3b indicate that the average number of bottles of treated wines sold increased slightly from March to April in the treatment store, while the average number of bottles sold decreased from March to April in the control store. Although this result suggests that the treatment increased consumer demand for treated wines, given the large standard errors, it is not possible to reject the null hypothesis that the change in the average number of wines sold is significantly different from zero. The final two columns in both tables indicate that most of the increase in the average number of treated wines sold in the treatment store was driven by increased demand for low-priced high-scoring wines. Specifically, the average number of low-priced high-scoring treated wines sold increased from 25.8 to 33.3 in the treatment store, while in the control store, the average number of such wines sold decreased from 23.5 to 20.5. The standard errors are, however, large and the difference in average quantity between months is not significantly different from

zero.

Results from the first columns of Tables 3a and 3b are supported by results from Specification 1, which are provided in Table 4. Although the coefficient for the treatment effect ($store * month$) is positive across all OLS specifications, it is never significant. Thus, the average effect of the treatment on the treated wines is not significantly different from zero. The only variable which is consistently significant is the promotion dummy. It is always positive, indicating that a wine placed on promotion 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 4 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.

Table 5 provides results for Specification 2. As in Table 4, the average effect of the treatment is positive across all specifications but is not significantly different from zero, and the promotion effect is positive and highly significant. The results also indicate that although there is no overall differential effect of quality on treated wine sales ($Store * Month * HighScore$), the treatment did affect the demand for low-priced high-scoring wines ($Store * Month * LowPrice * HighScore$). The coefficient is positive and significant across all specifications and indicates that the estimated effect on a treated low price wine of moving from low to high score lies between 8 and 15 additional bottles sold during a given month.¹³ Given that low-priced high-scoring wines sold an average of 26 bottles during March in the treatment store, sales increased by an average of 30 to 58 percent as a result of the treatment. Consumers therefore only responded to high quality information signals and only for a subset of wines, suggesting that not all publicity is good publicity. We briefly discuss two potential explanations for this finding. First, it may be that consumers who purchase expensive wines are more likely to be informed of wine quality and are

¹³It is important to note that due to data limitations, we are not able to control for the case in which wines become out-of-stock. However, it is likely that by not accounting for stock-outs, these estimates give a lower bound for the true effect of the treatment on low-priced high-scoring wines. That is, we expect that wines that were most likely to become out-of-stock were those wines which received high scores and that had relatively low prices. By not controlling for such stock-outs, our estimates are likely biased downwards.

therefore less likely to respond to additional quality information. Thus, the effect of expert opinion on consumer demand for expensive wines should be significantly reduced. Alternatively, it may take a long period of time before consumers fully adopt and trust a new source of information. When expert opinion information is posted, consumers may initially be skeptical of the information's accuracy. Instead of fully trusting the new information source immediately, consumers may attempt to verify the accuracy of the new source by sampling labeled wines for which the costs of experimentation are low. That is, they may only purchase wines that are relatively inexpensive. If the information provided by the expert opinion labels is then found by the consumer to be accurate, the individual may buy more expensive wines in the future. To the extent that this explanation is valid, it suggests that a treatment period of one month may not be sufficient to observe the full effect of expert opinion provision.

Tables 6 through 9 investigate the robustness of our results and show that the results presented in Tables 4 and 5 are likely not driven by unobserved time effects, and that they are robust to the use of different control stores and the use of the other treatment store. We first investigate the extent to which the results are sensitive to the choice of an alternate control store. Using the methodology described in Section 3, we select another control store with a residual mean that lies close to zero and present results from Specifications 1 and 2 in Tables 6 and 7. As the tables show, the average effect of the treatment on the treated is not significantly different from zero, while the effect of the treatment on low-priced high-scoring wines is consistently positive and significant. Further, the size of the estimated coefficient is comparable to that presented in Table 5.¹⁴

We next analyze the extent to which our results are sensitive to the use of the treatment store. Specifically, we now use the other treatment store and the residual methodology described in Section 3 to select an appropriate control. Results from Specification 2 are provided in Table 8.¹⁵ As in Tables 5 and 7, the average effect

¹⁴Although not presented, other control stores identified by the residual methodology also yield similar results.

¹⁵Results from Specification 1 are omitted for brevity, but as in Tables 4 and 6, the average treatment effect on the treated is not significantly different from zero.

on the treated is not significantly different from zero, however, the effect of the treatment on low-priced high-scoring wines is consistently significant and positive. The magnitude of the coefficient is now somewhat reduced, yet given that low-priced high-scoring wines sold an average of 8.6 bottles during March in the treatment store, sales of such wines increased by 47 to 106 percent as a result of the treatment. Although this increase is larger than that observed in the other treatment store, the range of increased sales overlaps for both stores. Therefore, overall, it appears that our results are not substantially affected by the treatment or control store used in the analysis.¹⁶

We next assign a false treatment to the March and April period in 2005 to examine the extent to which our results are generated by other external factors. For example, the effect identified in Tables 4 and 5 may be an artifact of seasonal or other advertising trends not observed in the data. Table 9 provides results for Specification 2 and shows that there is no significant and consistent effect of the treatment on low-priced high-scoring wines. In columns (1) and (2) the estimate is negative and significant, however, the coefficient estimate becomes insignificant once the promotion dummy is included. Although not reported, we assigned false treatments to all months between March 2005 and March 2006. In every case, the average treatment effect and the effect of the treatment on low-priced high-scoring wines is not significantly different from zero. Given that we find such an effect during the actual treatment month, and the significance and sign of the treatment effect is similar using other control and treatment stores, it appears that the treatment had no overall effect on treated wines, but that it did have a significant and positive effect on the demand for low-priced high-scoring wines.

Finally, we examine the impact of the treatment on untreated wines using Specification 3. As the results provided in Table 10 indicate, the treatment did not have a significant impact on untreated wines in the treatment store. Specifically, the coefficient on $Store*Month$ is generally marginally positive, but always insignificant. Thus, consumer demand for untreated wines apparently remained stable during the

¹⁶As with the discussion for Tables 6 and 7, the use of other control stores identified by the residual methodology with the alternate treatment store yield similar results.

treatment period in the treatment store. There are several potential explanations for this last finding. First, it may be that substitution by consumers towards treated wines and away from untreated wines was not one-for-one. That is, at least some consumers, when buying labeled wines, also continued to buy unlabeled wines. Alternatively, substitution may have been one-for-one, but consumers who previously did not purchase wine due to a lack of quality information entered into the wine market when expert opinion was provided. Finally, consumers may have substituted temporarily by stocking up on treated wines or spatially by reducing the quantity of wine purchased at competing stores.

6 Conclusions

Our results strongly suggest that expert opinion can affect the demand for wine by transmitting product quality information to consumers. 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. By randomly selecting 150 wines to display expert opinion information and through the selection of a control store with similar characteristics to those of the treatment store, we are 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 treatment 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. 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.

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. Rating agencies for wine, and for other products such as electronics, cars, and restaurants will thus likely allow consumers to make more fully informed 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.

References

- [1] Akerberg, D. 2001. “Empirically Distinguishing Informative and Prestige Effects of Advertising,” *The RAND Journal of Economics*, vol. 32, 316-333.
- [2] Akerberg, D. 2003. “Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Empirical Examination,” *International Economic Review*, vol. 44, 1007-1040.
- [3] Akerlof, G. 1970. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, vol. 84, 488-500.
- [4] Eliashberg, J. and S. Shugan. 1997. “Film Critics: Influencers or Predictors?,” *Journal of Marketing*, vol. 61, 68-78.
- [5] Jin, G. and P. Leslie. 2003. “The Effects of Information on Product Quality: Evidence from Restaurant Hygiene Cards,” *Quarterly Journal of Economics*, vol. 118, 409-451.
- [6] Meyer, B. 1995. “Natural and Quasi-Experiments in Economics,” *Journal of Business and Economic Statistics*, vol. 33, 151-161.
- [7] Montgomery, C. and B. Wernerfelt. 1992. “Risk Reduction and Umbrella Branding,” *Journal of Business*, vol. 65, 31-50.
- [8] Nelson, P. 1970. “Information and Consumer Behavior,” *Journal of Political Economy*, vol. 78, 311-329.
- [9] Reinstein, D. and C. Snyder. 2005. “The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics,” *Journal of Industrial Economics*, vol. 53, 27-51.
- [10] Sorensen, A. and S. Rasmussen. 2004. “Is Any Publicity Good Publicity? A Note on the Impact of Book Reviews,” Working Paper.

- [11] Wiggins, S. and W. Lane. 1983. "Quality Uncertainty, Search, and Advertising," *American Economic Review*, vol. 73, 881-894.
- [12] Wolinsky, A. 1995. "Competition in Markets for Credence Goods," *Journal of Institutional and Theoretical Economics*, vol. 151, 117-131.

Table 1: Store Characteristics¹⁷

	Treatment Store	Control Store	All Other Stores (Average)
Wine Sales Rank	23	31	
Wine Sales (2005 \$)	711,511	639,459	362,107
Number Bottles Sold	76,686	65,861	42,422
Percent Wine Sales of Total Grocery Sales	7.3	10.3	5.0
Percent Sales Premium Wine	8.7	10.9	5.0
Median Household Income (2005 \$)	115,299	129,274	72,134
Shelf Space (Linear Feet)	500	510	391

¹⁷This table provides descriptive statistics of store characteristics for the treatment store, the control store, and all other stores. The variables are defined in the following manner: (1) wine sales rank gives the number of stores above a given store that have higher wine sales in dollars, (2) wine sales and number bottles sold is measured for a 24 week period ending on 1/15/06, and (3) percent sales premium wines is the percent of sales during the 24 week period that were obtained from the sale of bottles with prices greater than \$8. Sales data are provided by Infoscan and median household income data are provided by the retailer.

Table 2: Descriptive Statistics for Treated and Untreated Wine in the Treatment Store¹⁸

	Treated Wines	Untreated Wines (With Scores)	Untreated Wines (Without Scores)
Score	84.1 [3.5]	83.7 [3.0]	
Quantity (March)	12.2 [20.3]	14.3 [19.9]	9.2 [18.2]
Quantity (April)	14.5 [21.9]	18.4 [20.0]	9.1 [18.0]
Price (March)	11.8 [7.8]	10.9 [6.3]	11.8 [9.0]
Price (April)	12.5 [10.3]	11.6 [7.2]	11.9 [8.9]
Percent Discounted (March)	57.1	64.0	54.2
Percent Discounted (April)	57.1	65.7	54.4
Percent Red	63.4	61.9	60.6
Number Observations	112	253	629

¹⁸For all continuous variables, we report the mean and standard deviation. Quantity gives the average number of bottles sold in a given month, price indicates the average price in a given month, percent discounted indicates the percentage of wines that were discounted in a given month, and percent red gives the percentage of wines that were red wines.

Table 3a: Descriptive Statistics for Treated Wines in the Treatment Store¹⁹

	All Wines	High Score Low Price	Non High Score Low Price
Score	84	86.8	83.5
	[3.5]	[1.0]	[3.6]
Quantity (March)	12.2	25.8	10.5
	[20.3]	[44.2]	[14.1]
Quantity (April)	13.5	33.3	11.3
	[21.9]	[48.6]	[15.7]
Price (March)	11.0	8.7	9.8
	[8.8]	[7.4]	[8.3]
Price (April)	10.7	9.5	9.8
	[10.8]	[7.2]	[10.5]
Percent Red	63.4	30.8	67.7
Percent Discounted (March)	57.1	76.9	54.6
Percent Discounted (April)	57.1	76.8	54.5
Number Observations	112	18	94

¹⁹The mean and standard deviation are provided for all continuous variables. A wine is considered to have a high score if its score is greater than or equal to an 86, and it is considered to be a low price wine if the per unit price is less than \$12. Quantity gives the average number of bottles sold in a given month, price indicates the average price in a given month, percent discounted indicates the percentage of wines that were discounted in a given month, and percent red gives the percentage of wines that were red wines.

Table 3b: Descriptive Statistics for Treated Wines in the Control Store²⁰

	All Wines	High Score Low Price	Non High Score Low Price
Score	83.9	86.4	83.6
	[3.5]	[0.7]	[3.6]
Quantity (March)	13.9	23.5	12.6
	[14.9]	[26.5]	[12.4]
Quantity (April)	13.6	20.5	12.6
	[14.4]	[24.1]	[12.4]
Price (March)	11.6	9.1	9.7
	[10.1]	[9.9]	[10.1]
Price (April)	11.7	8.8	10.1
	[10.2]	[10.0]	[9.8]
Percent Red	65.9	40.0	69.2
Percent Discounted (March)	72.7	90.1	70.5
Percent Discounted (April)	65.9	81.8	63.6
Number Observations	110	18	92

²⁰The mean and standard deviation are provided for all continuous variables. A wine is considered to have a high score if its score is greater than or equal to an 86, and it is considered to be a low price wine if the per unit price is less than \$12. Quantity gives the average number of bottles sold in a given month, price indicates the average price in a given month, percent discounted indicates the percentage of wines that were discounted in a given month, and percent red gives the percentage of wines that were red wines.

Table 4: OLS Results for Specification One²¹Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-1.62 [1.18]	-1.79 [1.20]	0.71 [1.35]	0.58 [1.37]	-0.07 [1.87]	-0.07 [1.84]
Treated Month	-0.25 [0.77]	-0.25 [0.77]	0.77 [0.95]	0.80 [0.97]	0.65 [1.11]	0.71 [1.23]
Store*Month	1.48 [1.07]	1.48 [1.07]	0.46 [1.36]	0.43 [1.39]	0.65 [1.43]	0.77 [1.55]
Red Dummy		-7.01* [4.22]		-7.93** [3.91]	-5.76 [3.62]	
Promotion Dummy			14.96*** [2.52]	15.43*** [2.62]	12.69*** [2.42]	13.04*** [2.54]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.01	0.03	0.15	0.19	0.20	0.22
Number Observations	400	400	400	400	400	400

²¹The regression is run using Specification 1 for treated wines in the treatment and control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level. Note that there are 400 observations since not all wines labeled in the treatment store were stocked at the control store during the treatment period.

Table 5 OLS Results for Specification Two²²Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-5.47**	-6.52**	-1.98	-3.21	-3.13	-3.24
	[2.11]	[2.38]	[1.94]	[2.06]	[2.83]	[2.89]
Treated Month	-3.51*	-3.97*	-0.09	-0.52	-0.67	-0.60
	[1.80]	[2.05]	[2.58]	[2.82]	[2.79]	[2.95]
Store*Month	3.10	4.29	0.34	1.84	2.54	2.40
	[2.47]	[2.74]	[2.46]	[2.74]	[2.68]	[2.77]
Low Price Dummy	5.36	3.73	4.17	1.81	2.75	
	[3.46]	[3.54]	[3.38]	[3.52]	[3.51]	
High Score Dummy	-4.68	-5.87*	-1.05	-2.50	-2.73	
	[3.21]	[3.36]	[3.06]	[3.23]	[3.18]	
Store*Month*Low Price	-2.52	-4.00	-1.11	-3.11	-3.55	-2.01
	[3.29]	[3.46]	[3.28]	[3.53]	[3.43]	[3.11]
Store*Month*High Score	-2.30	-3.86	1.66	-0.26	-1.06	-0.54
	[2.64]	[2.99]	[3.37]	[3.65]	[3.55]	[2.30]
Store*Month*Low Price*High Score	12.28**	14.51**	8.10*	10.96*	11.60*	8.22*
	[6.03]	[6.90]	[4.55]	[5.83]	[6.88]	[4.51]
Red Dummy		-4.17		-5.85*	-4.05	
		[3.34]		[3.24]	[3.20]	
Promotion Dummy			12.07***	12.85***	10.40***	8.72***
			[1.89]	[2.14]	[2.19]	[2.01]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.13	0.14	0.22	0.24	0.24	0.26
Number Observations	400	400	400	400	400	400

²²The regression is run using Specification 2 for treated wines in the treatment and control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.

Table 6: OLS Results for Specification One (Alternate Control Store)²³Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-1.87 [1.28]	-1.89 [1.29]	-1.25 [1.28]	-1.25 [1.30]	-1.60 [2.19]	-1.00 [2.16]
Treated Month	1.06 [0.75]	1.06 [0.75]	1.82* [1.04]	1.84* [1.08]	2.83 [1.69]	2.57 [1.68]
Store*Month	0.17 [1.03]	0.17 [1.03]	-0.58 [1.33]	-0.62 [1.37]	-0.63 [1.38]	-0.68 [1.53]
Red Dummy		-5.41 [4.56]		-7.22* [4.24]	-4.70 [4.40]	
Promotion Dummy			16.53*** [2.73]	17.25*** [2.93]	17.68*** [3.23]	15.52*** [3.04]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.01	0.02	0.15	0.17	0.18	0.19
Number Observations	444	444	444	444	444	444

²³The regression is run using Specification 1 for treated wines in the treatment and alternate control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.

Table 7: OLS Results for Specification Two (Alternate Control Store)²⁴
 Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-3.22	-3.57	-1.95	-2.57	-0.97	-1.23
	[2.42]	[2.72]	[1.61]	[1.99]	[2.95]	[2.73]
Treated Month	-2.94	-2.66	-1.44	-0.08	-0.92	-0.67
	[2.08]	[2.27]	[1.44]	[1.87]	[2.15]	[1.98]
Store*Month	2.33	2.64	1.77	3.08	2.65	2.89*
	[2.90]	[3.34]	[1.93]	[2.46]	[2.31]	[1.53]
Low Price Dummy	8.61**	8.13*	6.91*	6.46*	6.00*	
	[4.18]	[4.36]	[3.37]	[3.57]	[3.45]	
High Score Dummy	-2.99	-3.69	-1.29	-1.56	-2.02	
	[3.09]	[3.45]	[2.04]	[2.50]	[2.39]	
Store*Month*Low Price	-4.70	-5.03	-5.13*	-6.89	-6.25*	-6.64
	[4.15]	[4.62]	[2.89]	[3.48]	[3.18]	[3.37]
Store*Month*High Score	-1.58	-2.12	0.91	-0.85	-0.54	-1.02
	[3.08]	[3.55]	[2.53]	[3.07]	[2.89]	[3.05]
Store*Month*Low Price*High Score	12.34*	13.11*	9.60*	10.20*	11.83*	10.17*
	[6.76]	[7.26]	[5.74]	[5.69]	[6.27]	[5.51]
Red Dummy		-2.90		-5.24	-3.16	
		[3.67]		[3.52]	[3.96]	
Promotion Dummy			13.88***	15.28***	15.28***	16.04***
			[2.11]	[2.49]	[2.62]	[3.41]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.11	0.11	0.21	0.22	0.23	0.26
Number Observations	444	444	444	444	444	444

²⁴The regression is run using Specification 2 for treated wines in the treatment and alternate control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.

Table 8: OLS Results for Specification Two (Alternate Treatment and Control Store)²⁵
 Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-0.40 [2.72]	-0.19 [2.90]	0.23 [2.14]	0.89 [2.33]	5.16* [2.65]	4.02 [3.11]
Treated Month	-1.20 [2.43]	-1.11 [2.53]	0.04 [1.70]	0.38 [1.85]	-1.01 [2.24]	-0.83 [2.19]
Store*Month	0.02 [2.88]	-0.07 [2.99]	2.60 [2.06]	-2.97 [2.23]	-3.24 [2.10]	-2.41 [1.88]
Low Price Dummy	13.30*** [4.06]	13.30*** [4.09]	11.23*** [1.61]	11.16*** [3.39]	10.58*** [3.25]	
High Score Dummy	-2.94 [2.42]	-2.96 [2.47]	-1.29 [1.93]	-1.29 [1.98]	-0.85 [2.14]	
Store*Month*Low Price	-1.26 [4.02]	-1.15 [4.17]	1.99 [3.11]	2.45 [3.37]	3.25 [3.09]	4.00* [2.13]
Store*Month*High Score	-1.36 [3.04]	-1.35 [3.12]	2.67 [2.61]	2.85 [2.74]	3.45 [2.73]	2.62 [2.55]
Store*Month*Low Price*High Score	9.70* [5.88]	9.72* [5.98]	4.77* [2.51]	4.85* [2.60]	4.01* [2.31]	5.27* [2.68]
Red Dummy		-0.88 [2.98]		-2.70 [2.80]	-2.88 [3.56]	
Promotion Dummy			11.95*** [1.61]	12.33*** [1.73]	15.65*** [2.61]	9.81*** [2.03]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.15	0.15	0.26	0.26	0.27	0.28
Number Observations	444	444	444	444	444	444

²⁵The regression is run using Specification 2 for treated wines in the alternate treatment and control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.

Table 9: OLS False Treatment Results for Specification Two (March and April of 2005)²⁶
 Dependent Variable: Number of Bottles Sold of Treated Wine i in Store s during Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-1.17 [3.31]	-0.94 [3.45]	-2.19 [1.97]	-1.94 [2.10]	-0.39 [3.01]	-0.76 [2.88]
Treated Month	2.57 [4.10]	2.73 [4.27]	3.13 [2.44]	3.35 [2.64]	0.98 [3.63]	1.36 [4.02]
Store*Month	-3.65 [4.34]	-3.78 [4.55]	-2.81 [2.63]	-2.96 [2.84]	-3.01 [2.70]	-2.88* [1.52]
Low Price Dummy	8.18** [3.14]	7.26** [3.21]	2.34 [2.15]	1.03 [2.23]	0.23 [2.47]	
High Score Dummy	-0.58 [2.98]	-2.21 [3.33]	-2.75 [2.11]	-4.83* [2.45]	-5.37** [2.69]	
Store*Month*Low Price	5.31 [5.39]	5.36 [5.62]	0.72 [3.55]	0.64 [3.76]	0.63 [3.55]	0.83 [3.92]
Store*Month*High Score	7.05 [4.68]	7.39 [4.92]	0.57 [3.81]	0.79 [4.09]	0.80 [4.05]	1.06 [3.75]
Store*Month*Low Price*High Score	-19.90* [10.34]	-20.83* [10.19]	-6.64 [9.04]	-7.39 [9.23]	-6.76 [9.29]	-5.63 [11.14]
Red Dummy		-4.76 [3.18]		-5.87** [2.72]	-6.78** [2.92]	
Promotion Dummy			14.08*** [1.93]	14.50*** [1.94]	16.01*** [2.21]	15.57*** [2.10]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.12	0.13	0.28	0.31	0.31	0.33
Number Observations	350	350	350	350	350	350

²⁶The specification given for Model Two is run using data from the treatment and control store for the pre-treatment month and treatment month in 2005. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.

Table 10: OLS Results for Untreated Wines (Specification 3)²⁷
 Dependent Variable: Number of Bottles Sold of Untreated Wine i in Store s during
 Month t

	OLS (1)	OLS(2)	OLS (3)	OLS(4)	OLS (5)	OLS (6)
Treated Store	-1.92**	-1.94**	-0.59	-0.60	-1.19	-0.93
	[0.62]	[0.63]	[0.63]	[0.64]	[0.91]	[1.05]
Treated Month	-0.26	-0.26	0.39	0.40	0.23	0.18
	[0.41]	[0.41]	[0.51]	[0.51]	[0.80]	[0.47]
Store*Month	0.32	0.32	0.09	0.09	-0.42	-0.30
	[0.49]	[0.49]	[0.63]	[0.63]	[0.81]	[0.72]
Red Dummy		-7.70***		-8.34***	-8.15***	
		[1.90]		[1.78]	[2.07]	
Promotion Dummy			17.61***	17.91***	16.48***	15.55***
			[1.46]	[1.48]	[1.75] *[1.62]	
Low Price Dummy					3.94**	
					[1.84]	
Store*Month*Low Price					-0.63	-0.89
					[1.04]	[0.94]
Red & Promotion Interactions	No	No	No	No	Yes	Yes
Wine Fixed Effects	No	No	No	No	No	Yes
R^2	0.01	0.02	0.12	0.15	0.16	0.18
Number Observations	2,928	2,928	2,928	2,928	2,928	2,928

²⁷The regression is run using Specification 3 for untreated wines in the treatment and control store for the pre-treatment and treatment month. The red and promotional interactions are not significant and thus not reported. Standard errors are clustered by wine and are given in brackets. Standard errors are clustered by wine and are given in brackets. * Indicates that a point estimate is significant at a 10 percent level, ** indicates that a point estimate is significant at a 5 percent level and *** indicates that a point estimate is significant at the 1 percent level.