

NBER WORKING PAPER SERIES

CONSUMER BOYCOTTS:
THE IMPACT OF THE IRAQ WAR ON FRENCH WINE SALES IN THE U.S.

Larry Chavis
Phillip Leslie

Working Paper 11981
<http://www.nber.org/papers/w11981>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2006

Thanks to Dave Baron, Lanier Benkard, Paul Devereux, John McMillan, Garth Saloner, Andy Skrzypacz and Alan Sorensen for helpful advice. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Larry Chavis and Phillip Leslie. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Consumer Boycotts: The Impact of the Iraq War on French Wine Sales in the U.S.

Larry Chavis and Phillip Leslie

NBER Working Paper No. 11981

January 2006

JEL No. L0

ABSTRACT

The French Opposition to the war in Iraq in early 2003, prompted calls for a boycott of French wine in the US. We measure the magnitude of consumers' participation in the boycott, and look at basic evidence of who participates. Conservative estimates indicate that the boycott resulted in 26% lower weekly sales at its peak, and 13% lower sales over the six month period that we estimate the boycott lasted for. These findings suggest that business should be concerned that their actions may provoke a boycott which hurts their profits. We also find that neither political preferences or media attention are important determinants of boycott participation.

Larry Chavis

Stanford University

Graduate School of Business

518 Memorial Way

Stanford, CA 94305

chavis@larry@gsb.stanford.edu

Phillip Leslie

Stanford University

Graduate School of Business

518 Memorial Way

Stanford, CA 94305

and NBER

pleslie@stanford.edu

1 Introduction

Should businesses be concerned that their actions may provoke a consumer boycott which hurts their profits? Economists tend to be skeptical of the effectiveness of boycotts because of a free-rider problem. Individual consumers are glad for others to alter their purchase choices in support of some cause, but realize their own participation is unlikely to make any difference and would require some sacrifice. This intuition is reinforced by prior empirical studies into the effect of boycotts on firms' stock prices, which typically find small or insignificant effects. Nevertheless, purported boycotts are commonplace.¹ Recent examples include KFC (for alleged mistreatment of chickens), Nestlé (for marketing breast milk substitutes), Nike (for employment practices in East Asia), and Target (for not using the words "Merry Christmas" in its advertising). In this study we look at a direct measure of consumer participation in a boycott: the impact of a boycott on weekly product-level sales. Specifically, we measure the effect on sales of French wine from the US consumer boycott of French wine in 2003. We find a 13% decrease in sales over the six months we estimate the boycott lasted for. Hence, this example indicates that businesses should indeed be concerned about provoking a boycott of their products.

The French government did not support the US-led war in Iraq when it commenced on March 20, 2003. While France was not alone in their opposition to the war, as a permanent member of the United Nations Security Council, France was the most prominent of the opposing countries. Germany also opposed the war and was a temporary member of the security council at the time. However, France was more outspoken and more lambasted in the US-press. The first indication in a major US newspaper of a consumer boycott of French wine occurred in the *New York Times* on February 14, 2003. Of course the French wine industry played no role in the French government's opposition to the Iraq war. For consumers supporting the boycott of French wine, the hope is that somehow this may impact the behavior of the French government. Friedman (1999) defines this kind of boycott as a *surrogate boycott*, in which the French wine industry serves as a stand-in for the French government. Wine may not have been the only industry to experience a boycott of French products. Aside from boycotts, there were also other ways that people in America displayed their unhappiness with the French government, including attempts to rename *french fries* as *freedom fries*.

¹John and Klein (2003) argue that around 40% of *Fortune 50* companies may be subject to a boycott at any one time, and they note survey evidence indicating that 18% of Americans participate in boycotts.

In 2002, the year before the Iraq war, imports of French table wine accounted for 2.7% of the total volume of wine purchased in the U.S.² News reports describing the boycott of French wine in 2003 have provided conflicting indications as to whether there was any actual impact on French wine sales. Regardless, there are a couple of reasons to expect a non-trivial degree of consumer participation in the boycott. Firstly, there are numerous substitutes for French wine, and for at least some French wine, there are close substitutes from other countries. Hence, the sacrifice a consumer incurs by altering their purchase decision is likely to be minor. Secondly, the discontent towards France because of their opposition has been quite dramatic in the US. Gallup polls indicate that in May 2000, 50% of Americans considered France to be an ally and only 4% considered France to be unfriendly. However, in April 2003, 18% of Americans considered France to be an ally and 31% considered France to be unfriendly.

We obtained data for the period December 2001 to November 2003, in which we observe weekly price and quantity, by product and by city, for wine sales in mass-merchandise stores. The data cover four geographic markets in the US: Boston, Houston, Los Angeles and San Diego. We selected these cities because they have relatively high wine consumption per person, and because there is variation in political preferences—Boston and Los Angeles are Democrat-dominated regions, while Houston and San Diego are Republican-dominated regions. Importantly, for each wine product the data includes the country-of-origin. We identify the timing of the consumer boycott of French wine based on articles in leading national newspapers. Complete details of the data are provided in Section 2.

We focus on three main questions about consumer boycotts. First, how large was the effect of the boycott on French wine sales? Second, who participated in the boycott? Third, what impact did different types of media have on the magnitude of the boycott?

Our conservative estimate is that the boycott caused a 13% decrease in the volume of French wine sold over the first six months after the US war with Iraq commenced. In the conclusion we describe a back-of-the-envelope calculation indicating that total imports of French wine to the entire U.S. were lowered by \$112 million because of the boycott. The strength of the boycott varies from week to week. We estimate the peak of the boycott occurred nine weeks after the first news reports of the boycott, with an estimated 26% lower volume of French wine sold,

²*Adams Wine Handbook 2003*, p. 43. The revenue share of French wine would be significantly higher than 2.7% for 2002, due to the relatively high average price of French wine.

than if there had been no boycott. The strength of the boycott fades over time. Our estimates indicate that around six months after the boycott started, French wine sales are back to within 5% of where they would have been if there was no boycott. By the end of our sample, which is eight months after the war commenced, we find no significant impact from the boycott on weekly French wine sales.

Our finding that there is a significant degree of consumer participation in the boycott is at odds with the prediction of the simple free-rider theory. This is analogous to voting where economists also predict free-riding will lead to zero voter turnout, which is clearly wrong. The theory may be a straw-man that imposes unrealistic assumptions about individual rationality and self-interest. Nevertheless, the challenge is to better understand what drives individuals' participation decisions.

We examine three potential determinants of boycott participation. First, whether political preferences affect participation. The variation in Presidential voting across cities allows us to examine this aspect. We find the highest degree of participation in San Diego (Republican) followed by Los Angeles (Democratic) then Houston (Republican).³ Hence, the data indicates that participation is not closely aligned with political preferences. Second, whether willingness-to-pay for the boycott product affects participation. To do so, we estimate the impact of the boycott by French wine price-quartile. We find that cheap and expensive French wine are the most affected, while moderately-priced French wine is the least impacted. We conjecture that cheap wine buyers may have mild preferences for specific wines, and buyers of expensive French wine tend to give the wine to others as a gift. Hence, these buyers incur little disutility from substituting to wines from other regions. While buyers of moderately-priced wine have stronger preferences for French wine and intend consuming it themselves.

A third potential determinant of boycott participation that we consider is the role of the media. We focus on the importance of front-page coverage, and the outspoken support for the boycott by news media personality Bill O'Reilly of the *O'Reilly Factor* on Fox News. Our estimates suggest that front page news is no more impactful than non-front page news, and that Bill O'Reilly did not affect the magnitude of the boycott. Our prior belief was that political preferences and media attention would significantly affect boycott participation. Apparently the reason why people choose to participate in boycotts is more subtle.

³As we explain in Section 4 the data for Boston is unreliable.

Several prior papers analyze the impact of boycotts on the stock prices of target companies.⁴ Some find negative effects on stock prices: Friedman (1985), Pruitt and Friedman (1986), Pruitt, Wei and White (1988) and Davidson, Worrell and El-Jelly (1995). Other studies find no significant effect, or even positive effects: Koku, Akhigbe and Springer (1997) and Teoh, Welch and Wazzan (1999). The most recent paper of this kind, Epstein and Schnietz (2002), finds mixed evidence. We are aware of one previous study examining sales data for evidence of an effective boycott. Bentzen and Smith (2002) study aggregate monthly sales of French wine in Norway around the time of French nuclear testing in 1995-96, which prompted calls to boycott French products. Their analysis suggests there may have been a slight decrease in sales near the time of the nuclear tests, but does not quantify the effect or provide any statistical test of the claim. To the best of our knowledge, our study is the first to examine product-level data for evidence of boycott participation.⁵

In Section 2 we summarize the data. Section 3 contains our analysis of the effect of the boycott on aggregate French wine sales (i.e., quantity). In Section 4 we examine who participates in the boycott. The role of the media is analyzed in Section 5 and Section 6 is the conclusion.

2 Data Summary

There are two main components to our dataset: wine sales data and newspaper coverage of the French wine boycott. The sales data comes from Information Resources Inc (IRI) and is scanner data from supermarkets and other general merchandize stores. A limitation is that the data does not include sales at specialty wine stores or restaurants.⁶ However, a strength of the data is that it is weekly observations at the product level on a city-by-city basis, for the two year period of December 2001 to November 2003. Importantly, the data also identifies the country of origin of each wine product (or state if from the U.S.).⁷ All of the analysis in this study is

⁴A number of papers provide theoretical analyses of boycotts. See, for example: Baron (2003) and John and Klein (2003).

⁵Fershtman and Gandal (1998) use product-level data to measure the impact of the Arab boycott on Israel on consumer and producer welfare in the Israeli automobile market. In this case, Arab nations effectively stopped Japanese car manufacturers from selling products to Isreal. Consumer participation in the boycott was not an issue in that case.

⁶Off-premise sales of wine in 2002 for the entire U.S. accounted for 78.7% of all wine sales, by volume. See *Adams Wine Handbook 2003*, p. 30.

⁷We also observe the volume, name and type of wine for each product.

based on sales of 750ml bottles. The expense of the data limited us to obtaining it for four cities. We selected cities that vary in political preferences: Boston and Los Angeles are Democratic strongholds, and Houston and San Diego are Republican strongholds.

Table 1 summarizes the sales data based on the country or state of origin. There are 6,781 unique wines in the dataset, and 14,175 wine-city pairs. For these four cities, total wine sales (of 750ml bottles) over the two year period is over \$1 billion. Total wine sales for the entire U.S. in 2002 alone were about \$20.5 billion.⁸ Californian wines dominate our sample, with a 78.2% share of revenue. Wines from Italy are the second most common in the data, accounting for 6.3% of total revenue. French wines are equal third in revenue share with Australia (4% revenue shares for both). However, the average price of French wine is much higher than wines from any other region, making French wine the 5th most popular on the basis of unit shares, in these cities.

In Table 2 we compare the four cities in our data. French wine is relatively more popular in Boston with a 5% unit share, and the least popular in Los Angeles and San Diego. The two Californian cities exhibit a strong preference for wines from California. We had anticipated this when choosing the four cities, which is why we decided on two cities both from southern California with differing political preferences. We also report the average number of 750ml units per person in each of the cities. This measure varies considerably across the cities, from 0.44 in Boston to 8.73 in San Diego. Rather than revealing true differences in wine consumption, we take this as evidence that IRI's coverage of wine selling retailers is relatively poor in Boston and Houston, compared to Los Angeles and San Diego. This limitation of the data may impact our analysis. We assume the data is equally representative across the four cities, allowing us to make comparisons in the effectiveness of the boycott across cities. In Table 2 we also report the percent of votes for Bush (Republican) and Gore (Democrat) in the 2000 presidential election in each of the cities. It is apparent that Boston is strongly democratic, Los Angeles is democratic, San Diego is republican and Houston is strongly republican.

There is a question as to how to determine when the boycott is active or being called for. We implement two approaches in the analysis below. First, we define a French boycott dummy

⁸*Adams Wine Handbook 2003*, p. 8. The figure for total US sales includes table wine, wine coolers, champagne and sparkling wines, dessert and fortified wines, and vermouth/aperitifs. Table wine accounts for 90% of the aggregate, by volume. The total figure also covers wine in sizes other than 750ml bottles.

equal to one during the first eight weeks after the war commenced on March 20, 2003.⁹ This allows us to estimate straightforward difference-in-difference specifications as a basic indication of the effectiveness of the boycott. However, this approach ignores variation in the intensity of the boycott, and requires somewhat ad-hoc assumptions about when the boycott started and ended.

We therefore utilize a second approach based on newspaper reports that mention the words “France” or “French” in the headline and “boycott” in the text as a proxy for when the boycott is taking place. This has the appeal that we rely on a data source for determining when the boycott is active, as opposed to our judgement. Furthermore, the number of news articles in a given week is a measure of the intensity of the call for a boycott. We count articles in the leading national papers: *New York Times*, *Wall Street Journal* and *USA Today*. In addition, a high-profile proponent of the boycott was Bill O’Reilly on the Fox News channel. For some of the analysis we examine the separate effect of Bill O’Reilly, based on a count of the number of times he discussed the boycott each week.

For most of the analysis in this study we interpret the news variables as proxying the call for a boycott of French wine. Our primary goal is to assess the degree of consumer participation in the boycott, not whether the newspapers themselves had a causal impact on the boycott. Our interpretation is that the actual call for the boycott comes from a variety of sources, including politicians, media celebrities (such as Rush Limbaugh in addition to Bill O’Reilly) and other prominent individuals (such as Hollywood publicist Michael Levine). However, in reality the news coverage may be crucial for stimulating consumer participation, and so there may be some causal impact from the newspapers on the effectiveness of the boycott. Hence, in some portions of our analysis, as explained below, we explore the role of the media by separating the impact of front page news articles from non-front page articles.

Table 3 provides summary statistics of the news reports for the boycott. During the period of our data there were a total of 22 articles about the boycott in these three newspapers. Nine of these articles were on a front page. Bill O’Reilly discussed the boycott in 24 shows. Of the three newspapers, the *New York Times* had the most articles and the *Wall Street Journal* had the fewest. In the bottom panel of Table 3 we report the correlations of the various sources. It is comforting that all are positively correlated, which suggests the news articles may be a

⁹In fact the data is weekly, so this boycott period is defined as March 17, 2003 to May 11, 2003.

reasonable proxy for the boycott.

To better illustrate the data, in Figure 1 we plot weekly market shares for wines by country-of-origin, over the two year sample period, for the four cities in our dataset. We also include vertical bars (units on the right-side vertical axis) showing the weeks with newspaper reports of the boycott. We show the market shares for a set of four comparable countries—Australia, France, Italy and Spain. The diagram emphasizes the point that we observe sales for more than a year before the boycott, allowing us to identify underlying trends in sales for wines from each region. For example, it is clear from the first year of our data that Australian wine sales have a strong increasing trend. Also, while hardly conclusive, it is apparent that the French wine share falls at the time of the news reports about the French wine boycott. This is suggestive that there was some degree of participation in the boycott.

3 Effect of Boycott on French Wine Sales

In this section we measure the effect of the boycott on French wine sales. In the first subsection we estimate a difference-in-difference specification using weekly-product level observations for wines from all regions. In the second subsection we estimate a nonlinear model that allows us to measure the week-to-week variation in the strength of the boycott. The third subsection contains a robustness check where we estimate the effect of the boycott on wines from countries other than France. In subsequent sections we explore the mechanism of the boycott.

3.1 Difference-in-Difference Analysis

We begin our analysis of the effect of the boycott on French wine sales by examining the most basic evidence. In Figure 2 we show the percent change in the quantity of French wine sold each month in the second year of the data (November 2002 to October 2003) compared to the same month in the first year of the data. This is a simple way of controlling for seasonality.¹⁰ We also show the percent changes for wines from a selection of other regions, which provide useful comparisons. Importantly, the figure includes a vertical line showing the time of the first

¹⁰This also allows for the possibility that the seasonality in sales is different for wines from each region.

newspaper article mentioning the boycott. The figure is a graphical analog to the difference-in-difference specification we estimate below.

Figure 2 shows that French wine sales were lower in each month from November 2002 to September 2003 than compared to the same month a year before. This indicates that French wine sales were declining even before the boycott. In the final month of our sample period, October 2003, French wine sales are higher than one year before. Over the same period, wines from California, Italy and Spain have fluctuated between higher and lower sales than one year before. Australian wines are an exception, showing higher sales in every month, and sometimes to a dramatic degree—sales are more than 60% higher than a year before in three of the twelve months.

Focusing on the French wine time-series in Figure 2, the curve exhibits a mild u-shape. In particular, the change in French wine sales is the most negative during the three months after the start of the boycott. This suggests the boycott may have had a negative impact on French wine sales. However, sales of Californian, Italian and Spanish wines also drop markedly in March 2003, casting doubt on this interpretation. Overall, Figure 2 provides weak evidence, at best, of the boycott’s impact.

A straightforward method for estimating the impact of the boycott on French wine sales is to implement a difference-in-difference approach. Let Q_{ijkt} equal the quantity sold of wine i , in city j , originating from region k , in week t . We define the variable $Boycott_{kt}$ as a dummy variable equal to one for French wine during the two-month period March 17, 2003 to May 11, 2003 (the first eight weeks after the war commenced). We estimate the following specification:

$$\ln(Q_{ijkt}) = \alpha_{ij} + \tau_t + \theta \text{Boycott}_{kt} + \epsilon_{ijkt},$$

where α_{ij} are fixed-effects for each wine-city pair, τ_t are week fixed effects, θ is the coefficient of interest, and ϵ is the residual.

The inclusion of wine-city fixed-effects assures that identification of the boycott coefficient is based on within-wine-within-city variation in relative sales of French wine. The weekly time dummies τ_t control for general seasonality in wine sales. However, there may still be differences in seasonality for wines from different regions. To help limit any bias in the estimate of θ from idiosyncratic seasonality in French wine, we first estimate the above specification using data for

the two months the boycott variable is switched on, combined with the same period of time one year before. Hence, identification of θ is limited to variation in the relative sales of French wine during the same two month period one year beforehand, where it is plausible that seasonal demand for French wine relative to wines from other regions is constant. This is a sample of 226,800 wine-city-week observations. The estimate for θ is reported in first row of Table 4. With this specification, we estimate that the boycott caused an 8.8% decrease in French wine sales (significantly different from zero with 99% confidence).

Table 4 contains the results for an array of alternative specifications—each row is a separate regression. The main point to presenting these alternatives is to show that the negative effect of the boycott on French wine sales appears to be robust. Although the precise magnitude is variable. We now discuss each of the alternative specifications. For all the estimates in Table 4 we report robust standard errors. In every case, the estimate of the boycott coefficient is significantly different from zero with 99% confidence.

A concern may be that the estimate of the boycott effect in the first specification is due to stocking decisions of the retailers (or distributors) rather than consumers' choices. To address this, in the second specification we limit the sample to wine-city pairs for which we observe strictly positive sales in every week during the same four months as the first specification. We presume that these wines are always available for consumers, and hence any variation in sales is due to consumers' choices. This reduces the sample to 107,821 observations. As reported in the second row of Table 4, the estimate for the effect of the boycott is now larger—a 12.1% decrease in French wine sales.

We seek to interpret the estimate of the boycott coefficient as a demand response. However, if prices of French wines increased at the time of the boycott, this may also explain the reduction in French wine sales. Moreover, we could not rule out the possibility that prices of French wines are raised in response to the boycott, because high elasticity consumers may be more likely to participate in the boycott than low elasticity consumers. In the third specification we include the log of price on the right-hand side. Although not reported in the table, the estimated coefficient on $\ln(\text{Price})$ is -1.20 (standard error of .03). In this case the estimate for the boycott coefficient implies a 7.8% decrease in French wine sales. The negative coefficient on price, and the reduction in magnitude of the boycott effect relative to the second specification, indicate that relative prices of French wine may indeed have risen at the time of the boycott. This is

verified in the fourth specification where $\ln(\text{Price})$ is the dependent variable in a specification that is equivalent to the second row in Table 4. We estimate that the boycott caused a 1.1% increase in the price of French wine relative to wines from other regions.

While a price increase in response to the boycott is not out of the question, one would typically expect a reduction in demand to result in lower prices. Also, during the period of our data, the US dollar has been depreciating relative to the Euro which could explain rising U.S.-dollar prices of French wine. We therefore suspect that the estimate of the positive effect of the boycott on price may be spurious. To examine this possibility, in rows (5), (6) and (7) in Table 4, we report the results from re-estimating the prior specifications using only data on wines from European countries, for which the exchange rate effect is neutral. In this case we find effect of the boycott on prices is negligible (less than one percent decrease in relative price of French wine). The estimate for the impact of the boycott on French wine sales is now estimated to be around 16%.

Finally, a weakness of the above specifications is the absence of separate time-trends for wines from each region. If French wine sales have been trending down relative to sales of wines from other regions, as Figure 2 indicates, the above estimates will overstate the impact of the boycott. To address this concern, we use the full dataset, not just the four month samples used above, to estimate the following specification:

$$\ln(Q_{ijkt}) = \alpha_{ij} + \tau_t + \sum_k \left(\beta_{1k}t + \beta_{2k}t^2 + \beta_{3k}t^3 \right) + \theta \text{Boycott}_{kt} + \epsilon_{ijkt},$$

which includes origin-specific time trends (up to a cubic). The estimate for θ is reported in the final row of Table 4. In this case we find a 5.1% decrease in French wine sales due to the boycott.

3.2 Analysis of Weekly Boycott Intensity

The above difference-in-difference analysis indicates the boycott caused a decrease in French wine sales by an amount somewhere between 5.1% and 16.6%. To better gauge the magnitude of the effect we estimate a specification that allows the intensity of the boycott to vary from week to week. Also, rather than assume the boycott lasted for two months, as we did in the difference-in-difference specification, this approach yields an estimate of how long the boycott lasted.

Let Q_{kt} be the quantity of wine from region k purchased in week t (aggregated across all four cities in our data). We estimate the following model:

$$\ln(Q_{kt}) = \sum_k \left(\alpha_{0k} + \alpha_{1k}t + \alpha_{2k}t^2 + \alpha_{3k}t^3 \right) + \tau_t + \beta H_{kt} + \theta N_{kt} + \epsilon_{kt}, \quad (1)$$

where

$$N_{kt} = I_k^F (n_t + \delta N_{k,t-1}). \quad (2)$$

The variable n_t is the number of news articles in week t (in the New York Times, Wall Street Journal and USA Today) with the words “France” or “French” in the headline and “boycott” in the text. I_k^F is an indicator variable equal to one for France and zero otherwise. Hence, N_{kt} measures the intensity of the boycott—it is the depreciated stock of boycott news articles. The model also includes region fixed-effects (α_{0k}), region-specific time trends and week fixed-effects. Also, because empirically French wine is particularly popular on certain holidays, we include a holiday dummy H_{kt} which equals one for French wine in weeks with a major holiday.¹¹

There are two key parameters of interest. Firstly, δ measures the rate of depreciation of participation in the boycott. If $\delta = 0$, calls for a boycott last week have no impact on boycott activity this week. We expect that $0 < \delta < 1$.¹² The closer that δ is to one, the longer the boycott lasts. Secondly, θ measures the contemporaneous response of consumers to current calls for boycotting French wine. The more consumers that participate in the boycott, the more negative will be θ . With estimates of δ and θ in hand, and data on news articles (n_t), we can compute the variable impact of the boycott in each week. We estimate the model via nonlinear least squares.

The results are reported in Table 5, including estimates for differing sub-samples and variations on the above specification. The top row of Table 5 is based on the full sample (4,160 observations). We obtain very precise estimates of both δ and θ : $\hat{\delta} = 0.86$ and $\hat{\theta} = -0.02$. Based on the high R^2 values shown in the table, we conclude that the model provides a good fit to the data. A potential concern with this analysis is serial correlation in the dependent variable. However, we compute a Durbin-Watson statistic of 1.95 for French wine, indicating the absence of any significant serial correlation. Also, in Figure 3 we show actual and predicted sales of French wine, where it is apparent that we provide close predictions in almost all periods. The figure also includes counterfactual sales, which we explain below.

¹¹The specific holidays are Valentine’s Day, Thanksgiving, Christmas, and New Years.

¹²This is not a constraint imposed for estimation.

The estimates themselves are not very intuitive measures of the boycott’s magnitude. Hence, we compute three other measures of the implied magnitude of the boycott, as shown in the last three columns of the table. In each case, we compare predicted sales of French wine given the boycott, with the predicted sales if there was no boycott. To compute the counterfactual we set $n_t = 0$ in all periods, and compute predicted quantities based on the estimated parameters. The time-series of the counterfactual is shown in Figure 3, where the counterfactual shows higher sales from around February to July, 2003.

Comparing factual and counterfactual sales in each week, the first of the three measures is the *maximum weekly effect*. We find that, at the peak of the boycott, weekly French wine sales would have been 26.6% higher if there was no boycott. A second measure of the boycott’s magnitude is the percent of lost sales over the six months following the start of the boycott (February 10, 2003 to August 17, 2003). Again, the calculation is based on the counterfactual described above. For the base specification, we find that French wine sales were 13.3% lower than what they would have been if there was no boycott, over this six month period. A third measure of the boycott’s magnitude is the estimated duration, defined as the number of months until French wine sales return to within 5% of what they would have been if there was no boycott. In the top row of Table 5 we report the estimated duration to be 5.7 months, for the base specification.¹³

Figure 4 graphically depicts the estimated weekly variation in boycott intensity, based on the counterfactual described above. We also include vertical bars showing the timing and quantity of news articles referring to the boycott. The time path of the boycott magnitude reflects the instantaneous responses to boycott calls, followed by periods of depreciation in the degree of participation. While the magnitude is above 25% at only one point, there are 18 consecutive weeks where the reduction in sales due to the boycott exceeds 10%.

The estimates reported in the remaining three rows of Table 5 serve as robustness checks. By almost any measure, the alternative specifications we consider give rise to larger boycott effects. In the second row, we include $\ln(\text{Price})$ as an independent variable. Since an observation is the aggregate quantity of wine for a given region-of-origin in a given week, price is defined as the weighted average price. We now only include wines with positive sales in a given week, and so the number of observations falls to 2,808. As shown in Table 5, we find a larger maximum

¹³Note the dataset extends about 9 months after the start of the boycott.

weekly effect of the boycott (now 39.8% versus 26.6% under the base model). The six month effect and duration are also larger than the base model.

As discussed above in the difference-in-difference analysis, it may be reasonable to limit the sample to only European wines. In the third row of Table 5 we report the results of using this sample for the nonlinear model. The implied magnitudes are quite similar to the base model. Finally, we estimate the nonlinear model using only the data for French wine, so that identification comes entirely from the time-series of French wine sales. With two years of weekly data, this implies 104 observations. As shown in Table 5, the estimates for θ and δ are still very precise. The three measures for the magnitude of the boycott in this case indicate the largest of all—maximum weekly effect of 45%, 26% lower sales over six months, and the boycott duration of 8.5 months.

Hence, the estimates for the above nonlinear specification suggest a conservative estimate is that the French wine boycott lasted around 6 months, during which time French wine sales were approximately 13% lower, and at the peak of the boycott weekly sales were down by about 26%.

3.3 Effect of Boycott on Sales of Non-French Wines

We expect the boycott of French wine causes some degree of substitution to wines from other countries. Hence, it would provide verification of the boycott effect if we found the boycott caused an increase in sales of wines from other regions. This also serves as a specification check. If we find that the boycott causes lower sales of wines from other regions, this would suggest our empirical analysis does a poor job of separating the effects of the boycott from underlying time-trends.

In Table 6 we present results for the difference-in-difference model as well as the nonlinear model, in which we estimate the effect of the French wine boycott on sales of wines from other regions. Each row contains the results from two different specifications. In the top row we restate the results for the effect of the boycott on French wine sales as a comparison.

Looking at the impact on Australian wine sales, the difference-in-difference model indicates

3.8% higher sales due to the boycott. This seems reasonable in comparison to the estimated 12.1% decrease in French wine sales, based on the same model. Recall, for the difference-in-difference model we consider the effect of the boycott over a two month period. The nonlinear model indicates 12.1% higher sales over the six months following the start of the boycott. This is probably too high—some of the high growth in Australian wine sales may be incorrectly attributed to the boycott.

For sales of Californian wine we obtain unrealistic results under both specifications. First, the difference-in-difference model indicates the boycott caused lower sales of Californian wine. In contrast, the nonlinear model implies almost 19% increases in sales of Californian wine. Given the large market share of Californian wine, the estimated increase is many times larger than the decrease in sales of French wine. These results cast doubt on our analysis more generally. The fact that Californian wines are so dominant in the US, with around 80% market share, suggests it may be unreasonable to compare sales of French wine with sales of Californian wine.

The results for wines from Chile and Italy are more encouraging. The signs and magnitudes are reasonable under both models. Since Chile has a smaller market share than France, the six month estimate of 13.4% increased sales is not implausible. Finally, the results for Spain are mixed. The difference-in-difference estimate is reasonable, but the nonlinear model implies an implausible effect, as shown in the bottom row of Table 6. On the one hand, the results shown in Table 6 provide mixed evidence as to how convincing our estimates are of the impact of the boycott on French wine sales. On the other hand, these results do suggest that wines from Australia, Chile and Italy have benefitted from the French wine boycott.

Finally, in an unreported regression, we estimate the effect of the boycott on sales of Californian wines with French-sounding names, such as the winery *Chateau Julien*. We found no significant effect on sales for these wines. This could be due to consumers' ability to recognize such wines as being non-French, or because stores tend to shelve wines by country-of-origin which helps consumers avoid confusion.

4 Who Participates in the Boycott?

Who participated in the French wine boycott? We consider two characteristics of potential participants. Firstly, are Republican supporters more likely to boycott French wine than Democrat supporters? Gallup polling suggests that Republicans may be more likely to boycott French wine: in February 2004, 64% of Republicans and 37% of Democrats held an unfavorable opinion towards France.¹⁴ The second characteristic we consider is whether buyers of cheap or expensive French wine more likely to participate in the boycott.

We do not observe consumer-level decisions on whether to boycott French wine. However, we observe product-level sales for each geographic market, varying in aggregate political preferences. Hence, we estimate the effect of the boycott by price-quartile, and we estimate the effect of the boycott separately for each of the four cities in our dataset.

We start by analyzing the variation in the boycott effect by political preferences. Figure 5 shows the time-series of the market share for French wine in each city. As noted above, Boston and Los Angeles are pro-Democrat markets, and Houston and San Diego are pro-Republican markets. The figure reveals that the data for Boston provides a very different pattern of sales than for the other three markets. The increase in market share for French wine in Boston in 2003 may be due to a reverse-boycott effect in Boston, sometimes referred to as *buycott*, or may simply be due to a problem with the data. Since the increase in Boston begins a couple of months **ahead** of the boycott, it seems much more likely to indicate a substantive mid-sample change in IRI's data collection in Boston, than anything related to the boycott. Hence, while we present the estimates for Boston, these results should not be taken seriously.¹⁵

For the remaining three cities in Figure 5, we see the now familiar pattern of lower market shares for French wine corresponding to the boycott news articles. Los Angeles and San Diego offer a nice comparison, since the time-series are very similar for both. It also appears that the boycott was longer lasting in Los Angeles than in San Diego. To quantify the boycott effect for each market, we separately estimate the nonlinear model shown in equations (1) and (2) for each

¹⁴By comparison, in February 2002, prior to the war in Iraq, 15% of Republicans and 16% of Democrats held unfavorable views of France. See "Image of France Begins to Recover in American Eyes", *The Gallup Organization*, February 18, 2004.

¹⁵Also, please note that we re-estimated all the previous results in the paper with the data for Boston excluded, and found no qualitative differences in any of the findings.

city. The results are presented in Table 7. All coefficient estimates are significantly different from zero with 99% confidence.

We focus on the implied measures of the boycott magnitude and ignore the results for Boston. As shown in the table, the estimates imply San Diego has the largest maximum weekly effect and the largest 6-month effect of the boycott. Both of these measures show that Los Angeles has the second largest effects, and Houston is third. The ranking is reversed based on the duration measure, which may be right, or else suggests our method is poor at separately identifying the θ and δ coefficients. Of the three measures, the 6-month effect is arguably the most relevant, because it reflects a combination of the intensity and duration of the boycott.

On face value, our estimates suggest the boycott was most effective in San Diego, followed by Los Angeles. Hence, it does not appear that boycott participation is closely aligned with political preferences. The reason why the boycott was strongest in San Diego may be due to the strong military presence (Navy and Marines) in the area, but this is speculation.

We now examine the effectiveness of the boycott for wines in different price categories. Was participation in the boycott greater for the buyers of cheap or expensive French wine? Figure 6 shows the time-series of the market share of French wine in each quartile of the distribution of French wine prices. Note that the average price of French wine is well above that of wines from other regions. Hence, in the top price quartile (prices above \$25.76) French wines dominate the market, with around 40% share in that category. To make the figure more readable, we divide the share of French wine in the top price quartile by a factor of 10.

In the figure, the French wine share in the lowest price quartile (prices below \$7.98) noticeably falls at the time of the boycott. In the top price quartile there is distinct downward spike in market share of French wine during the boycott, although it appears to be transitory. For the middle quartiles, the time-series reveal no obvious evidence of boycott effects.

To quantify the effects of the boycott, we again re-estimate the nonlinear model given by equations (1) and (2) separately for each price quartile (aggregating across cities). The results are presented in Table 8. All coefficient estimates are significantly different from zero with 99% confidence. As Figure 6 indicated, the boycott of expensive French wines was intense but short-lived. We compute a maximum weekly effect of 52% lower sales, but a 6-month effect of only a

12.4% decrease. It is noteworthy that the impact on expensive wines was large but short-lived, and it is an appealing feature of our specification that we are able to separately identify these two dimensions of boycott participation.

Cheap French wine, on the other hand, has a maximum weekly loss of 73.8% and a 6-month loss of 34.4%. The impact on medium-to-low priced French wine (priced between roughly \$8 and \$12) is similar to the cheap wines, but not quite as severe. The medium-to-expensive French wines (around \$12 to \$26) are the least impacted, and in fact our estimates imply the boycott caused small increase in sales for these wines.

We can speculate on the explanation for why the strength of the boycott differs across price categories in this way. The reason why expensive French wine has a big impact from the boycott may be because people are more likely to give such wine as gifts (relative to other wine).¹⁶ The reasoning is that the buyer of the wine does not intend to consume it themselves, so there is little disutility from substituting to wine from another region. Also, the gift-giver may feel a desire to show the gift-receiver that they are participating in the boycott.

The reason why cheap French wine has a big impact from the boycott may be because consumers of cheap wine tend to be indifferent between different wines, which is why they consume cheap wine. Again, for these consumers the disutility from substituting to wine from another region may be low. Conversely, the buyers of moderately-priced French wine are a group of consumers that may have more refined tastes for wines, and they purchase this wine for their own consumption rather than as gifts. Hence, these consumers are the least likely to be willing to switch to wines from another region. Of course these explanations are entirely speculative.

To summarize the results of this section, we find that: (i) participation in the French wine boycott does not seem to be related to political preferences; and (ii) buyers of cheap French wine and expensive French wine seem to be the most likely to participate in the boycott, with medium-priced French wines the least impacted.

¹⁶In Figure 6, high price French wine is the only category to display a dramatic spike in sales in the holiday period.

5 Role of the Media

The above findings show that there was an economically significant degree of participation in the French wine boycott. Thus, skepticism that people would not participate due to the free-riding problem appears to be unfounded. But can we generalize from these results to other boycotts? Two features of the French wine boycott may heighten consumers' participation beyond what we may expect in other examples. Firstly, the availability of close substitutes. Although this is certainly not unique to wine. Secondly, the French wine boycott received a fair amount of media attention. In this section, we attempt to shed light on the role of the media in stimulating participation in the French wine boycott. Specifically, we examine whether front page news reports have a bigger impact on the boycott than non-front page articles. Also, since Bill O'Reilly of the *O'Reilly Factor* on Fox News has been a high-profile proponent of the boycott, we examine whether his comments on the boycott have increased participation.

To test the hypotheses that front page news coverage is an important determinant of boycott participation, or that advocacy by a prominent media personality such as Bill O'Reilly is important for stimulating participation, we re-estimate the nonlinear model given in equations (1) and (2) with different coefficients for each media source. It is important to highlight the change in assumptions we now make. Previously, we assumed that news reports about the boycott are a proxy for the underlying calls for a boycott. Hence, we interpreted the estimated effects as simply measuring consumer participation, rather than the causal effect of news reports on participation. In contrast, in this section we seek to identify the causal impact of different news sources on boycott participation. To do so, we separately include multiple news sources in the specification, and test the distinct effects of each. In doing so, we estimate which particular news source is most highly correlated with participation in the boycott.

This empirical strategy is more compelling for identifying the causal effect of Bill O'Reilly than the effect of front page news. The reason is because news coverage of the boycott is more likely to be on the front page when the underlying call for the boycott is strongest. We therefore expect that front page news appears to have bigger impact than non-front page news, even if there is no causal effect. However, it is conceivable that the tendency of Bill O'Reilly to discuss the boycott on any given day is random. Indeed, we note that during the 6-month period of May 2005 to October 2005, long after the boycott started, Bill O'Reilly has continued to mention the

French boycott on at least 9 occasions.

Table 9 reports the estimates for four versions of the nonlinear model with multiple news sources. The first column of estimates is for the base model previously discussed, to provide a comparison. The second column of estimates is for a model that includes only the weekly count of shows in which Bill O'Reilly mentions the boycott. Recall, the estimate for θ captures the instantaneous impact and the estimate for δ captures the longevity of the effect. The significance and size of the second column of estimates suggests Bill O'Reilly may have been an important determinant of boycott activity.

However, when newspaper articles and the O'Reilly mentions are separately included, as shown in the third column of estimates in Table 9, it does not appear that Bill O'Reilly stimulated participation in the boycott.¹⁷ Indeed, the estimate of the instantaneous impact of Bill O'Reilly is for an increase in sales of French wine. Perhaps Bill O'Reilly stirs a backlash against the boycott. Regardless, it does not appear that his advocacy was an important driver of participation.

In the final column of estimates in Table 9 we examine the separate effects of front page and non-front page news. As we discussed, there is good reason to doubt that we have identified the causal effects of these news sources. Nevertheless, the estimates suggest that both types of news articles have about the same effects, both in terms of the instantaneous effect ($\hat{\theta}$) and the longevity ($\hat{\delta}$). On face value, this may imply that front page news coverage is not particularly important for promoting boycott participation.

¹⁷Recall, in Table 3 we report the correlations between news coverage of the boycott from different sources, including Bill O'Reilly. The correlation between O'Reilly and the various newspapers is positive but never greater than 0.5.

6 Conclusion

By examining weekly product-level sales data, we find that there was economically significant consumer participation in the boycott of French wine in 2003 in the United States. Alternative specifications, as well as the use of various subsamples, all indicate a non-trivial degree of boycott participation. However, the precise magnitude of the effect is less clear. Our preferred specification, the so-called nonlinear model, has the appeal that it relies on news articles for determining when the boycott is active, and allows us to separately identify the intensity and longevity of consumers' participation. With this model, we conservatively estimate that, at the peak of the boycott, the quantity of French wine sold would have been 26% higher if there was no boycott. Also, over the six month period that we estimate the boycott lasted for, sales would have been 13% higher. But some estimates indicate a maximum weekly effect of more than 40% forgone sales, and 20% lost sales over six months.

How large is the absolute value of lost revenue? To calculate this, we compute the mean price of French wine conditional on each quartile of the distribution of French wine prices, then multiply by the estimate for total quantity of French wine sold in each price quartile over the six month period of the boycott, and sum together. This procedure gives the prediction for French wine revenue of \$7.4 million, for Boston, Houston, Los Angeles and San Diego.¹⁸ If we recompute the predicted quantity sold in each price quartile under the assumption that there was no boycott, and calculate total revenue, we obtain \$8.6 million. Hence, our estimates imply lost revenue of \$1.2 million, for these four cities, which is equal to 16.1% of the observed revenue for French wine over this six month period.

From March 2003 to August 2003, roughly the time of the boycott, the total value of wine imported to the U.S. from France was equal to \$695 million.¹⁹ If we extrapolate our revenue calculation and assume this value would have been 16.1% higher had there been no boycott, then the revenue loss for French wine importers to the U.S. market during this time was approximately \$112 million. Of course this is a crude back of the envelope calculation.

Our findings show that, despite the free-riding problem economic theory suggests would in-

¹⁸Actually, the exact prediction is \$7,409,385. Note that the observed revenue is \$7,409,541, which is remarkably close, suggesting this is a reasonable approach.

¹⁹The exact figure is \$694,822,551, and is defined as the "landed duty-paid value" of wine imported from France to the U.S. for the period March 2003 to August 2003, from the U.S. International Trade Commission.

hibit participation, consumer boycotts can be an effective approach for pressuring organizations to change their behavior. John and Klein (2003) propose a number of possible psychological explanations for why the free-riding problem may not stifle boycott participation. For example, individual consumers may have an exaggerated sense of their own effectiveness, or individuals may have a false sense of consensus. Suffice to say, empirical testing of alternative explanations for boycott participation is beyond the scope of this paper.

Lastly, the results of this study also verify that firms' foreign-earned profits can be harmed by their government's foreign policy. Whether these kind of economic effects are taken into account by governments when setting foreign policy is a question for future research. Or put differently, do economic incentives impose limits on the kinds of foreign policies that can be implemented? For example, it has been suggested that governments may be unlikely to adopt foreign policies that are critical of an important trading partner. Consistent with this view, our analysis sheds light on one specific mechanism via which business profitability depends on foreign policy.

Bibliography

- Baron, D.P. (2003): "Private Politics," *Journal of Economics and Management Strategy*, 12(1), 31–66.
- Bentzen, J. and V. Smith (2002): "French Nuclear Testings and Consumption of French Wine in Denmark," *International Journal of Wine Marketing*, 14(3), 20–36.
- Davidson, W.N., D.L. Worrell and A. El-Jelly (1995): "Influencing Managers to Change Unpopular Corporate Behavior Through Boycotts and Divestitures," *Business & Strategy*, 34, 171–196.
- Epstein, M. and K. Schnietz (2002): "Measuring the Cost of Environmental and Labor Protests to Globalization: An Event Study of the Failed 1999 Seattle WTO Talks," *The International Trade Journal*, 16(2), 129–160.
- Fershtman, C. and N. Gandal (1998): "The Effect of the Arab Boycott on Israel: The Automobile Market," *RAND Journal of Economics*, 29(1), 193–214.
- Friedman, M. (1985): "Consumer Boycotts in the United States, 1970–1980: Contemporary Events in Historical Perspective," *Journal of Consumer Affairs*, 19, 96–117.
- Friedman, M. (1999): *Consumer Boycotts: Effecting Change Through the Marketplace and the Media*, Routledge.

- John, A. and J.G. Klein (2003): "The Boycott Puzzle: Consumer Motivations for Purchase Sacrifice," *Management Science*, 49, 1196–1209.
- Kaempfer, W.H. and A.D. Lowenberg (1999): "Unilateral Versus Multilateral International Sanctions: A Public Choice Perspective," *International Studies Quarterly*, 43, 37–58.
- Klein, J.G., N.C. Smith and A. John (2004): "Why We Boycott: Consumer Motivations for Boycott Participation," *Journal of Marketing*, 68(3), 92–109.
- Koku, P.S., A. Akhigbe, and T.M. Springer (1997): "The Financial Impact of Boycotts and Threats of Boycotts," *Journal of Business Research*, 40, 15-20.
- Pruitt, S.W. and M. Friedman (1986): "Determining the Effectiveness of Consumer Boycotts: A Stock Price Analysis of Their Impact on Corporate Targets," *Journal of Consumer Policy*, 9, 375–387.
- Pruitt, S.W., K.C.J. Wei and R.E. White (1988): "The Impact of Union-Sponsored Boycotts on the Stock Prices of Target Firms," *Journal of Labor Research*, 9, 285–289.
- Teoh, S.H., I. Welch and C.P. Wazzan (1999): "The Effect of Socially Activist Investment Policies on the Financial Markets: Evidence from the South African Boycott," *Journal of Business*, 72, 35–89.

Table 1: Market summary by origin of wine for sales in Boston, Houston, Los Angeles and San Diego, over the period November 2001 to October 2003.

	Revenue (\$)	Revenue Share (%)	Quantity	Quantity Share (%)	Mean Price (\$)	Number of Products
California	859,585,857	78.2	102,668,966	78.6	8.37	7,593
Italy	69,635,676	6.3	7,852,725	6.0	8.87	1,325
France	44,369,842	4.0	2,841,079	2.2	15.62	1,415
Australia	43,927,773	4.0	5,161,468	4.0	8.51	1,065
Washington	21,807,524	2.0	2,289,560	1.8	9.52	318
New York	17,468,967	1.6	4,301,908	3.3	4.06	205
Chile	12,364,580	1.1	1,781,645	1.4	6.94	523
Spain	10,953,638	1.0	1,174,432	0.9	9.33	317
Texas	5,678,569	0.5	822,676	0.6	6.90	133
Germany	2,366,678	0.2	394,296	0.3	6.00	144
Other	11,699,820	1.1	1,339,130	1.0	8.74	1,137
TOTAL	1,099,858,923		130,627,884		8.42	14,175

Table 2: Overview of city characteristics

	Boston	Houston	Los Angeles	San Diego
	Percent of total units			
California	58	61	82	82
Italy	11	7	5	5
France	5	3	2	2
Australia	14	10	3	4
Total quantity	2,344,982	13,861,788	80,735,444	24,773,377
2002 population	5,309,000	4,713,500	15,752,400	2,837,500
Units per person	0.44	2.94	5.13	8.73
Vote for Bush in 2000	32%	57%	41%	50%
Vote for Gore in 2000	60%	40%	55%	46%

Table 3: News coverage of French wine boycott

	New York Times	USA Today	Wall Street Journal	Bill O'Reilly on Fox	Total
Number of news items					
All stories	13	6	3	24	46
Front page	6	3	0	NA	9
Correlation between news sources					
New York Times	1				
USA Today	.08	1			
Wall St Journal	.41	.20	1		
Bill O'Reilly	.29	.44	.45	1	

Table 4: Difference-in-difference analysis of boycott effects

Dependent variable	Boycott coefficient	Standard error	Sample	$\ln(\text{Price})$ on RHS	Origin-specific time trends	Obs	R^2
(1) $\ln(\text{Quantity})$	-.0880	0.0098	4 months	No	No	226,800	.86
(2) $\ln(\text{Quantity})$	-.1213	0.0173	4 months, sales>0	No	No	107,821	.89
(3) $\ln(\text{Quantity})$	-.1120	0.0173	4 months, sales>0	Yes	No	107,821	.90
(4) $\ln(\text{Price})$	0.0113	0.0023	4 months, sales>0	No	No	107,821	.98
(5) $\ln(\text{Quantity})$	-.1613	0.0207	4 months, sales>0, Euro-only	No	No	21,644	.88
(6) $\ln(\text{Quantity})$	-.1665	0.0207	4 months, sales>0, Euro-only	Yes	No	21,644	.88
(7) $\ln(\text{Price})$	-.0066	0.0029	4 months, sales>0, Euro-only	No	No	21,644	.98
(8) $\ln(\text{Quantity})$	-.0511	0.0056	All data	No	Yes	1,474,200	.83

All regressions include wine-city fixed effects and week fixed effects. Robust standard errors are reported.

Table 5: Estimates of boycott effect for nonlinear specification

Sample	$\hat{\theta}$	Standard error	$\hat{\delta}$	Standard error	$\ln(\textit{Price})$ on RHS	Obs	R^2	Max weekly effect	6-month effect	Duration (months)
(1) Full sample	-0.0238	0.0053	0.8595	0.0468	No	4,160	.99	-26.6%	-13.3%	5.7
(2) Full sample	-0.0333	0.0030	0.8647	0.0130	Yes	2,808	.99	-39.8%	-19.9%	6.5
(3) Euro only	-0.0303	0.0001	0.7899	0.0012	No	416	.99	-27.4%	-11.9%	5.7
(4) France only	-0.0313	0.0002	0.9157	0.0005	No	104	.77	-45.1%	-26.2%	8.5

All regressions include holiday dummies for France, origin-specific time trends, and week fixed-effects (specification (4) does not include week fixed-effects). The 6-month effect is based on the period 2/10/03 to 8/17/03. Duration estimates are defined as the number of months from the start of the boycott until sales of French wine are within 5% of the level they would have been if there was no boycott. Robust standard errors are reported.

Table 6: Boycott effect for regions other than France

	Difference-in-Difference		Nonlinear Specification				
	Coefficient	Std error	$\hat{\theta}$	Std error	$\hat{\delta}$	Std error	6-month effect
France	-.1213	0.0163	-.0238	0.0053	0.8595	0.0468	-13.3%
Australia	0.0380	0.0303	0.0285	0.0061	0.8254	0.0498	12.1%
California	-.0428	0.0155	0.0119	0.0028	1.0089	0.0343	18.8%
Chile	0.0179	0.0386	0.0174	0.0041	0.9228	0.0389	13.4%
Italy	0.0213	0.0265	0.0025	0.0027	1.0066	0.1722	4.1%
Spain	0.0287	0.0386	-.0292	0.0036	1.0427	0.0033	-122.4%

All regressions are based on the full sample. For all estimates based on the difference-in-difference specification, we also include wine-city fixed effects and week effects. For all estimates based on the nonlinear specification, we include region-specific time trends, week fixed effects, and a France-specific holiday dummy.

Table 7: Effect of boycott on French wines sales in each city

Sample	$\hat{\theta}$	Standard error	$\hat{\delta}$	Standard error	Obs	R^2	Max weekly effect	6-month effect	Duration (months)
(1) Boston	-0.0378	0.0023	1.0344	0.0027	2,704	.99	-920.2%	-177.0%	≥ 8.5
(2) Houston	-0.0150	0.0029	0.9455	0.0225	3,120	.99	-22.3%	-15.5%	≥ 8.5
(3) Los Angeles	-0.0334	0.0100	0.8854	0.0582	2,704	.99	-43.2%	-23.7%	7.8
(4) San Diego	-0.0998	0.0097	0.7729	0.0135	2,288	.99	-115.5%	-40.4%	6.5

As noted in the text, the estimates for Boston are probably due to data problems and should not be taken seriously. All regressions include holiday dummies for France, origin-specific time trends, and week fixed-effects. The 6-month effect is based on the period 2/10/03 to 8/17/03. Duration estimates are defined as the number of months from the start of the boycott until sales of French wine are within 5% of the level they would have been if there was no boycott. Robust standard errors are reported.

Table 8: Effect of boycott on French wine sales by French price quartile

Sample	$\hat{\theta}$	Standard error	$\hat{\delta}$	Standard error	Obs	R^2	Max weekly effect	6-month effect	Duration (months)
(1) 0 – \$7.98	-.0551	0.0072	0.8639	0.0281	3,328	.99	-73.8%	-34.4%	7.8
(2) \$7.99 – \$11.93	-.0216	0.0011	0.9727	0.0084	2,912	.99	-41.5%	-30.5%	≥ 8.5
(3) \$11.94 – \$25.76	0.0060	0.0030	0.9588	0.0475	2,392	.99	8.4%	6.3%	0
(4) \geq \$25.77	-.1162	0.0106	0.1858	0.0841	1,560	.98	-52.0%	-12.4%	5.0

All regressions include holiday dummies for France, origin-specific time trends, and week fixed-effects. The 6-month effect is based on the period 2/10/03 to 8/17/03. Duration estimates are defined as the number of months from the start of the boycott until sales of French wine are within 5% of the level they would have been if there was no boycott. Robust standard errors are reported.

Table 9: Effect of media on French wine boycott

Media source		Coefficient estimate	Standard error						
Newspaper articles	$\hat{\theta}$	-0.238	0.0053						
	$\hat{\delta}$	0.8595	0.0468	-0.206	0.0040	0.9202	0.0598		
Bill O'Reilly	$\hat{\theta}$			-0.340	0.0068	0.0137	0.0080		
	$\hat{\delta}$			0.7744	0.0691	1.0398	0.0207		
Front page news	$\hat{\theta}$							-0.257	0.0117
	$\hat{\delta}$							0.8748	0.1726
Non-front page news	$\hat{\theta}$							-0.222	0.0134
	$\hat{\delta}$							0.8387	0.3974
Observations		4,160		4,160		4,160		4,160	
R^2		.99		.99		.99		.99	

All regressions include holiday dummies for France, origin-specific time trends, and week fixed-effects. Robust standard errors are reported.

Figure 1: Weekly wine market shares by region and boycott news articles

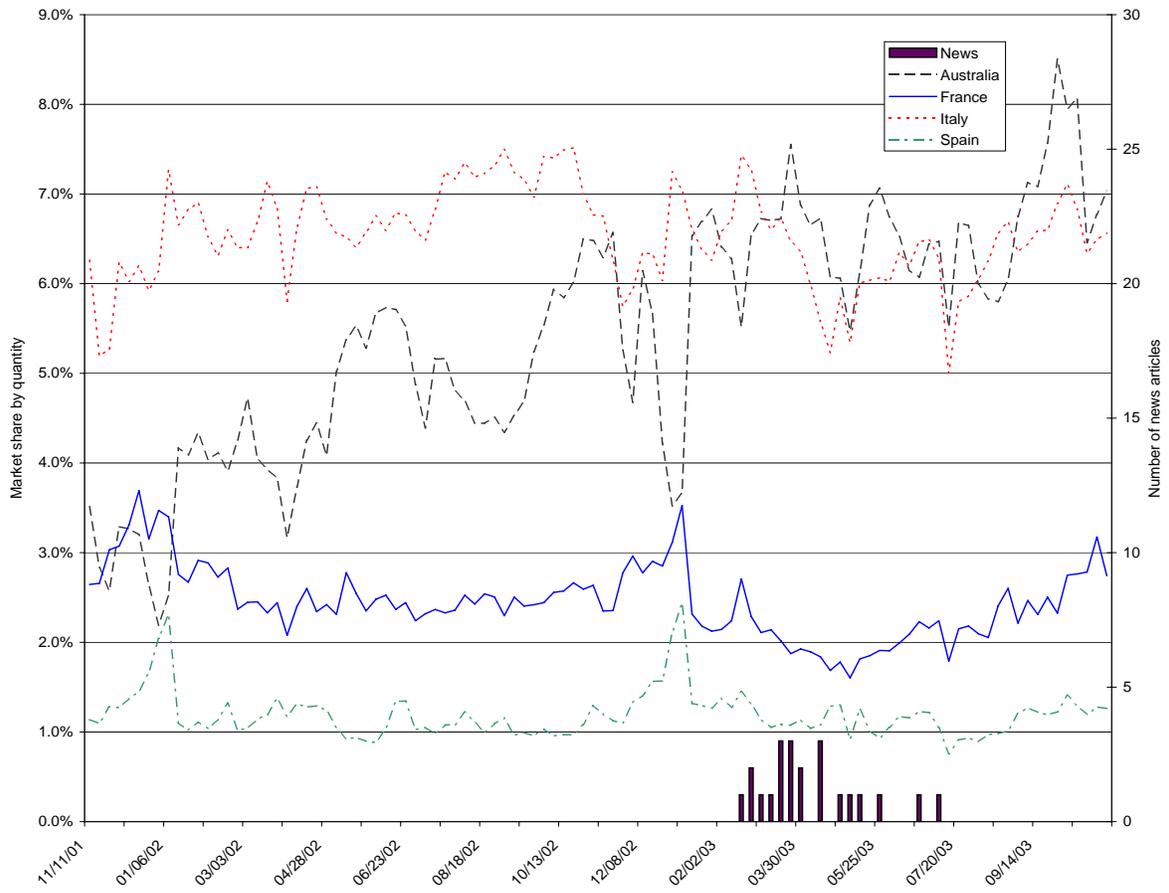


Figure 2: Month-on-month percent differences in wine sales by region-of-origin

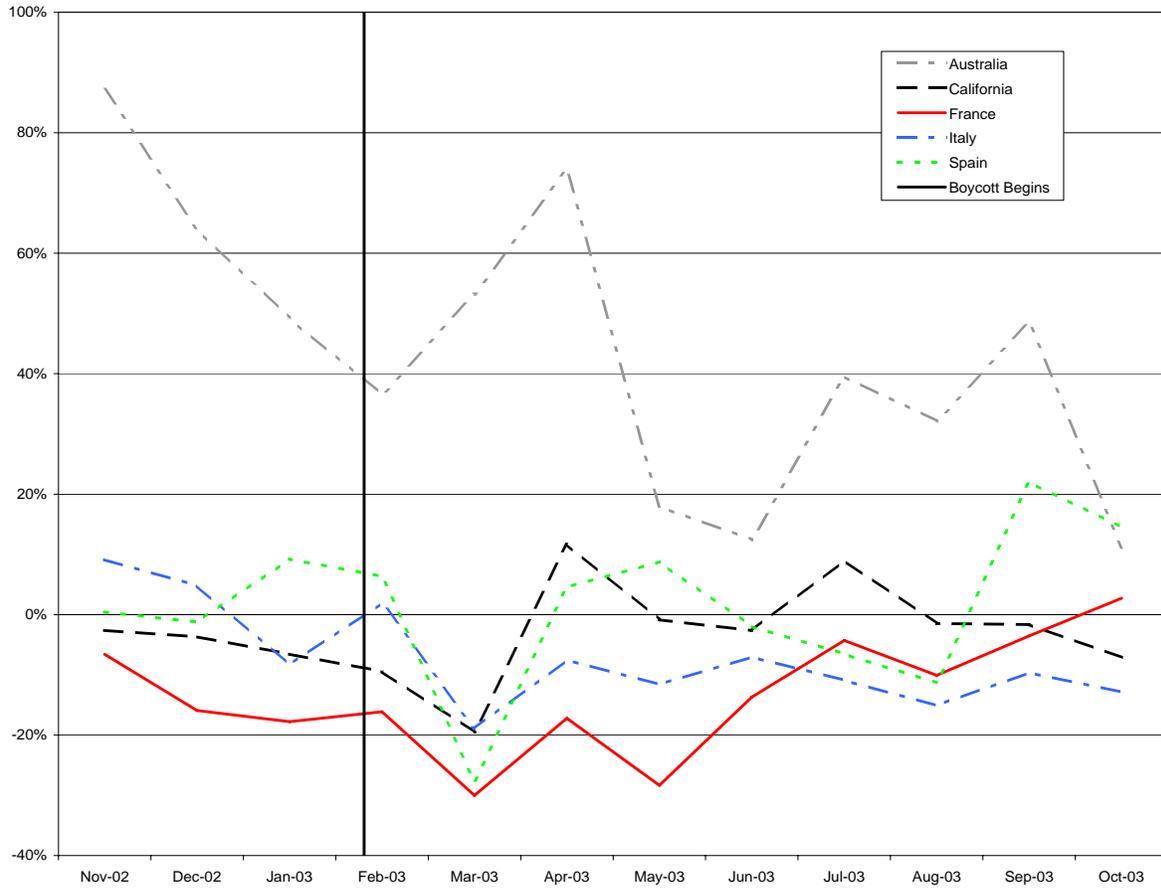


Figure 3: Actual, predicted and counterfactual French wine sales

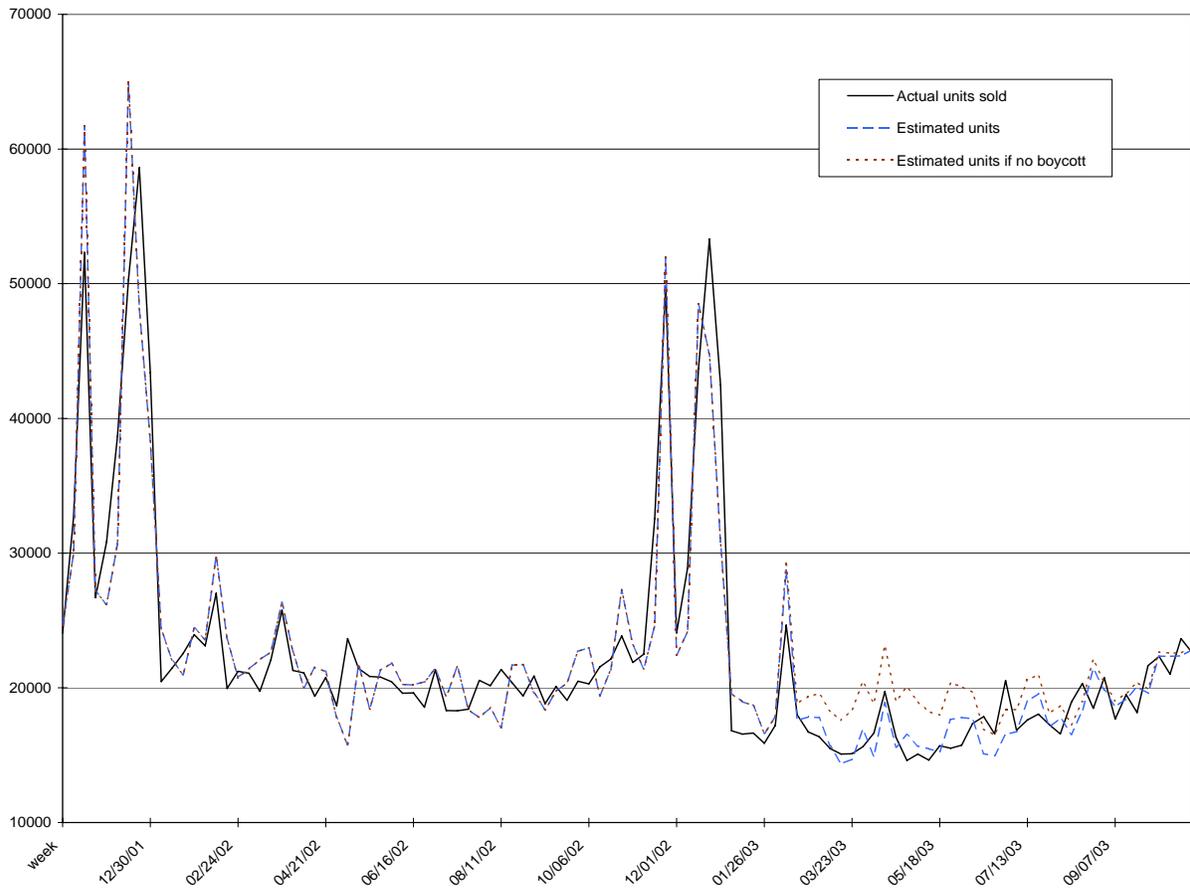


Figure 4: Estimated percent increase in French wine sales if there was no boycott

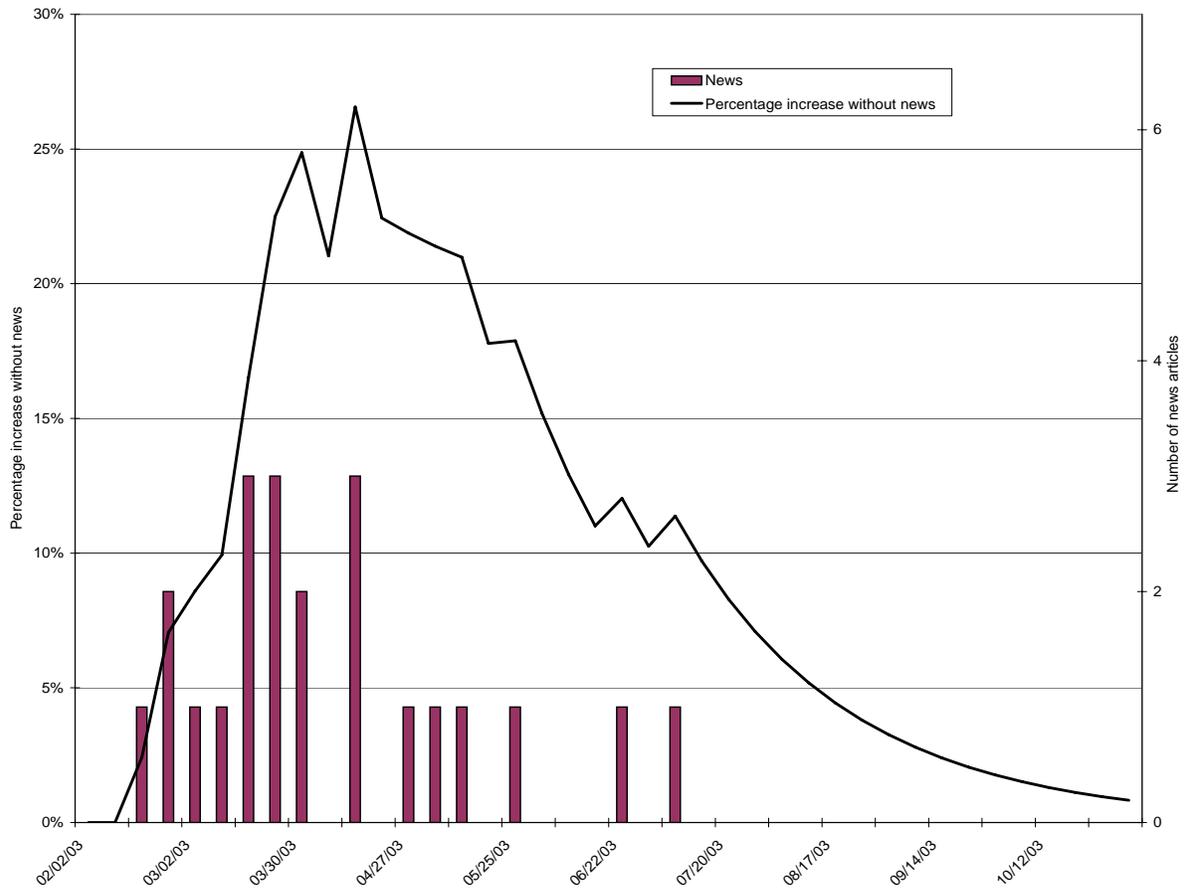


Figure 5: French wine market share by city and boycott news articles

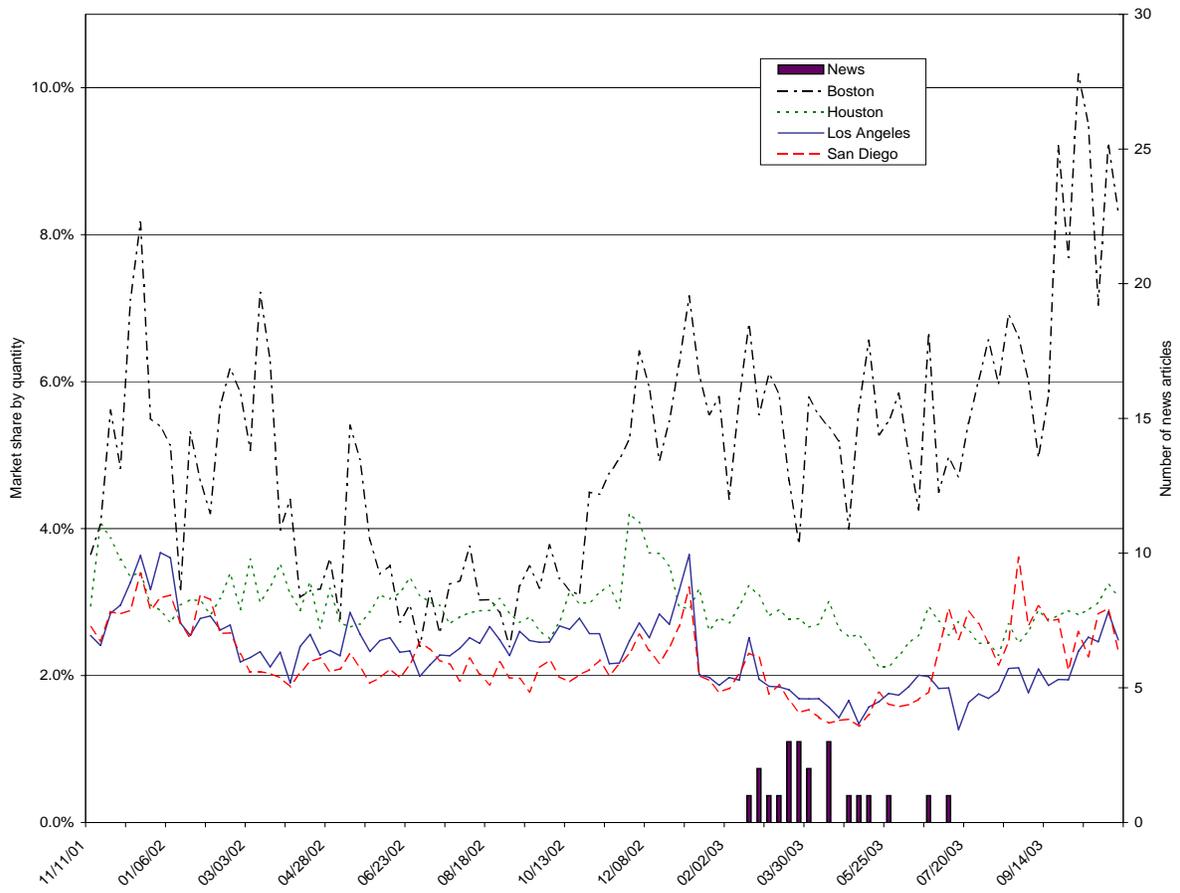
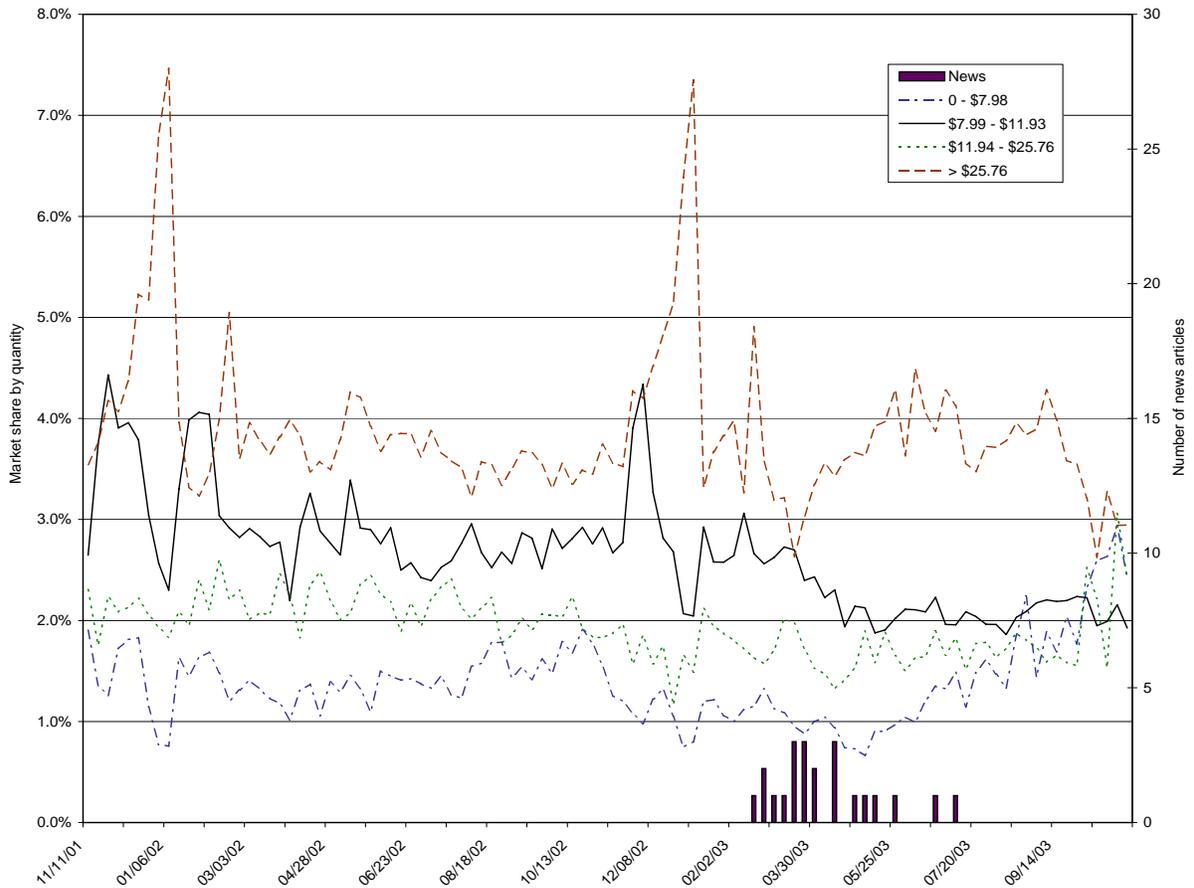


Figure 6: French wine market share by price quartile and boycott news articles



The market share of French wine in the top price quartile (price > \$25.76) is divided by 10 to make a clear figure. In fact, French wine dominates sales in the high price category.