

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

Larry Chavis  
Kenan-Flagler Business School  
University of North Carolina  
larry\_chavis@unc.edu

Phillip Leslie  
NBER, and  
Graduate School of Business  
Stanford University  
pleslie@stanford.edu

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. At the peak of the boycott, our conservative estimate is that French wine sales would have been 27% higher if there had been no boycott. We also find that the boycott lasted for six months, and French wine sales would have been 13% higher over these six months if there had been no boycott. Although theory suggests consumers would not participate in boycotts due to a free-rider problem, these findings indicate that businesses should be concerned that consumers may boycott their products. We also find that neither political preferences or media attention are important determinants of boycott participation.

July, 2008

---

\*Thanks to Dave Baron, Lanier Benkard, Paul Devereux, John McMillan, Harikesh Nair, Paul Oyer, Garth Saloner, Andy Skrzypacz and Alan Sorensen for helpful advice. We also thank our editor (Peter Rossi) and anonymous referees for valuable feedback.

# 1 Introduction

Calls for voluntary consumer boycotts of particular firms' products are commonplace. 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).<sup>1</sup> Given all this activity, firms should seemingly be very attentive to the threat of a potential boycott and being able to respond in ways that limit the harm to their profits. Or should they? Prior studies measuring the impact of boycotts on firms' stock prices find small or negligible effects. Moreover, a free-riding logic suggests consumers are unlikely to voluntarily participate: 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.

To examine this puzzle—whether managers should really be concerned about voluntary consumer boycotts—rather than look at indirect evidence (stock prices) we look at weekly product-level sales data. 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. Hence, this example indicates that businesses should indeed be concerned about consumer boycotts. The use of micro-level sales data has other benefits. We are able to examine the lifecycle of a boycott, regional variation in boycott participation, and boycott variation by price segment. We also explore the role of the media in stimulating boycott participation. Each of these factors has implications for how managers may respond to a boycott.

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

---

<sup>1</sup>Simply search the term “boycott” at Google to see the numerous current examples of purported boycotts. Or see the long list of current boycotts at EthicalConsumer.org. John 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.

opposition to the Iraq war. For consumers supporting the boycott of French wine, the hope was 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. There were other ways that people in America displayed their unhappiness with the French government, including attempts to rename *french fries* as *freedom fries*.

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.<sup>2</sup> 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 some degree of consumer participation in the boycott. Firstly, there are probably close substitutes for at least some French wine, making it easier for consumers to switch products. 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, and how did the intensity vary over time? Second, who participated in the boycott? Third, what impact did different types of media have on the magnitude of the boycott?

---

<sup>2</sup>*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.

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 27% lower volume of French wine sold, 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.

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).<sup>3</sup> 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. Hence, these buyers incur little disutility from substituting to wines from other regions. We also conjecture that buyers of expensive French wine tend to give the wine to others as a gift.<sup>4</sup> This has two consequences: the buyer of the wine is less likely to also consume the wine (making it more substitutable), and gift giving is an opportunity to make a public political statement.

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.

---

<sup>3</sup>As we explain in Section 4 the data for Boston is unreliable.

<sup>4</sup>In fact, in the data we observe that sales of high priced wine (and French wine in particular) dramatically spikes upward around Christmas time.

These findings have several implications for managers. Boycotts do in fact have the potential to significantly reduce sales for a firm, and so managers may wish to avoid certain actions that could prompt one. If there is a call for a boycott, the degree of participation may vary across geographic markets and across products (especially by price segment). Micro-level sales data can help to identify what kinds of customers are most actively participating in a boycott, which can guide a targeted response, such as price reductions or advertising campaigns, to mitigate the negative impact on profits. Collecting survey data would be an expensive alternative to obtaining similar information. The lifecycle of the French wine boycott exhibits a ramping-up period for two months, followed by gradual decay over six months. Hence, it can be important for managers to respond very early to a boycott, and to consider the possibility of lower sales for six months or more afterwards.

Several prior papers analyze the impact of boycotts on the stock prices of target companies.<sup>5</sup> 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.<sup>6</sup>

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.

---

<sup>5</sup>A number of papers provide theoretical analyses of boycotts. See, for example: Baron (2003) and John and Klein (2003).

<sup>6</sup>Fershtman 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 Israel. Consumer participation in the boycott was not an issue in that case.

## 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.<sup>7</sup> 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.).<sup>8</sup> All of the analysis in this study is 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 were about \$20.5 billion.<sup>9</sup> 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 and Australian wines each have 4% of revenue. 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 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

---

<sup>7</sup>Off-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.

<sup>8</sup>We also observe the volume, name and type of wine for each product.

<sup>9</sup>*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.

of the data may impact our analysis. 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 being called for. We implement two approaches in the analysis below. First, we define a French boycott dummy equal to one during the first eight weeks after the war commenced on March 20, 2003.<sup>10</sup> 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 to identify the period of time during which the boycott is being called. This has the appeal that we rely on a data source for determining when the boycott is being called for, 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*. The tone of these reports is almost entirely neutral, with a focus on describing the boycott and French anxiety over its effects.

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) 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. In addition, a high-profile proponent of the boycott was Bill O’Reilly on the Fox News channel. We also examine whether Bill O’Reilly had a causal effect on boycott

---

<sup>10</sup>In fact the data is weekly, so this boycott period is defined as March 17, 2003 to May 11, 2003. As a robustness check, we also try both longer and shorter time periods for the boycott.

participation, based on a count of the number of times he discussed the boycott each week.

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 reasonable proxy for the boycott call.

To better illustrate the data, in Figure 1 we plot weekly market share of French wine sales over the two year sample period, aggregated over the four cities in our dataset (unweighted). We also include vertical bars (units on the right-side vertical axis) showing the weeks with newspaper reports of the boycott. 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. The figure shows several positive demand spikes, including two around the beginning of 2003. The spikes correspond to Christmas, Valentines' Day and Thanksgiving. 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, although the changes may not be stark enough to merit strong evidence. Regression analysis allows us to explicitly control for country-specific time trends, helping to provide a clearer analysis of the impact of 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. The second subsection contains a robustness check where we estimate the effect of the boycott on wines from countries other than France. In the third subsection we examine evidence concerning the role that retailers may have had in the boycott. The fourth subsection we estimate a nonlinear model that allows us to measure the week-to-week variation in the strength of the boycott. In subsequent sections we explore the mechanism of the boycott.

### 3.1 Difference-in-Difference Analysis

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:

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

where  $\alpha_{ij}$  are fixed-effects for each wine-city pair,  $\tau_t$  are week fixed effects,  $\theta$  is the coefficient of interest, and  $\epsilon$  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  $\tau_t$  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  $\theta$  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. This allows us to control for seasonal variation in the demand for French wine relative to wines from elsewhere. In other words,  $\theta$  is identified from variation in sales of French wine relative to wines from other regions, holding fixed seasonal preferences. This is a sample of 226,800 wine-city-week observations. The estimate for  $\theta$  is reported in first row of Table 4. With this specification, the estimated coefficient on the boycott dummy is -.09, implying that the boycott caused an 8.4% 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. 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.

An important concern is whether the estimate of the boycott effect is due to a demand response by consumers, or a supply response by retailers or distributors. One way to address

this is to re-estimate the specification using only data for wine-city combinations that have strictly positive sales in all periods. For these observations it is more likely that bottles are available on the shelf at all points in time.<sup>11</sup> Hence, if we find decreased sales for these wines, it is likely to be due to a demand response rather than a supply response. 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 11.4% decrease in French wine sales (based on an estimated coefficient of -.12). Below, we examine in more detail whether retailers may have also changed their marketing mix because of the boycott.

Another concern relates to prices. If prices of French wines increased at the time of the boycott, this may also explain the reduction in French wine sales. Indeed, 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. Alternatively, prices of French wines may decrease at the time of the boycott, as an optimal response by retailers to lowered demand. In this case, ignoring the price effect would lead us to understate the degree of consumers' participation in the boycott, because price declines would stimulate sales.

There is an argument for ignoring price changes in this context: these price changes are also driven by the boycott, and the fact that a price change may exacerbate or mitigate the impact on French wine sales is as much a result of the boycott as consumers voluntary participation. On the other hand, a concern may be that retailers raise French wine prices due to their own desire to boycott French wine, rather than an optimal response to changes in demand, leading us to overstate the degree to which consumers choose to participate in the boycott. For this reason, it is important to control for price changes. To be clear, it is not our goal to estimate a price elasticity, which would warrant a more careful consideration of the sources of price variation in the data.

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 (-.11) implies a 10.6% decrease in French wine

---

<sup>11</sup>A caveat is that it only takes one store in a city to stock a wine and have positive sales for it to remain in this sample. Hence, a sales reduction may still be driven by other retailers removing wines from shelves. Note that we explore the role of retailers in more detail below.

sales. The negative coefficient on price, and the slight 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 1.0% 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 (because low elasticity consumers are less likely than high elasticity consumers to participate in the boycott), 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 14.9% (based on an estimated coefficient of -.1613).

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 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 with origin-specific time trends (up to a cubic):

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

where  $I_k$  is an indicator variable equal to one for country  $k$  and zero otherwise. The estimate for  $\theta$  is reported in the final row of Table 4. In this case we find a 5.0% decrease in French wine sales due to the boycott.

We also examined two other important robustness checks that are unreported in the tables. First, we vary the definition of the boycott period to be either longer (3 months: 2/10/03 to 5/11/03) or shorter (1 months: 3/17/03 to 4/13/03) than the 2 month window we use

above.<sup>12</sup> For the longer window there is no noticeable change in the estimates, and for the shorter window the estimates indicate even stronger boycott participation. Second, we estimate weighted specifications in which the weights are given by the total quantity sold for each wine-city pair. For every specification shown in Table 4, the estimated boycott effects become larger when the observations are weighted.

### 3.2 Effect of Boycott on Sales of Non-French Wines

A compelling test for whether consumers participated in a boycott of French wines is to simultaneously estimate the effect of the boycott on sales of wines from other regions. 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 find that the boycott caused an increase in sales of wines from other regions. But even more importantly, a finding that the boycott appears to lower sales from other countries would cast doubt on our conclusion that the boycott lowered sales of French wine.

To implement this test, we modify the difference-in-difference model above, by adding interactions between the boycott dummy and region dummies for Italy and Spain. We choose Italy and Spain because these countries were both supporters of the Iraq War, and because they share a common currency with France. In other words, we can rule out boycott participation extending to wines from these countries, and exchange-rate fluctuations are common with France.

The result from this “placebo-regression” is shown in Table 5. We report the results for both unweighted and weighted regressions (where the weights are the total sales in each wine-city pair). And we report the results when German wines are also interacted with the boycott dummy. Since Germany was also against the war, it is conceivable German wines were also boycotted, although none of the news coverage of the French boycott used in this study includes any mention of German products being targeted. To estimate these effects we use the sample of wines for which we observe strictly positive sales in all periods. This helps to insure that we measure demand-side responses.

---

<sup>12</sup>For the 3 month window we start the boycott period at the date of the first news article that mentions the boycott. For the 1 month window, and for the 2 month window in the base specifications, we start the boycott period at the beginning of the Iraq war.

As shown in Table 5, the estimated effect of the boycott on French wine sales is larger for the weighted-regression than in the unweighted-regression. This is because highly popular French wines are impacted more by the boycott than less popular wines. In the analysis in Section 4 we explore how this relates to the prices of French wines. The first column of estimates in Table 5 indicate that Italy experienced a statistically significant 4% increase in sales, and there was no significant effect on Spanish wines. In this specification, the impact on French wine remains significantly negative at around -11%. In the weighted regression, shown in the second set of estimates, the negative impact on France substantially increases, and the estimates for Italy and Spain are insignificantly different from zero. These findings support our conclusion that there was in fact significant consumer participation in the French wine boycott.

In the final two sets of results shown in Table 5 we also estimate the effect of the boycott on German wine sales. In the unweighted regression we find an implausible 10% increase in German wine sales. But when the observations are weighted, the estimated effect is statistically insignificant.

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.

### **3.3 Retailers' Role in the Boycott**

It is conceivable that retailers change their behavior in response to the boycott, in ways that may either enhance or mitigate the impact of the boycott. For example, retailers may support the boycott by reducing shelf-space of French wine (beyond any reduction that may be an optimal response to lower demand). Or retailers may increase their promotion activity of French wine to reduce the impact of the boycott on French wine sales. In this subsection we explore the possibility that the boycott caused a change in retailers' French wine marketing mix. It would be ideal to utilize data on how much shelf space is allocated by retailers to wines from different countries, or to have data on advertising, neither of which is available to us. However, we examine a few measures that are likely to be correlated with retailers marketing mix choices

more generally.

We observe the regular price, promotional price, quantity sold at the regular price and quantity sold at promotional price for each product-city-week observation in the data. Elsewhere in this study we use the average price and total quantity sold (i.e. regular plus promotional) in our analysis. In this subsection we examine the disaggregated measures for evidence that retailers may have changed the marketing mix. In Table 6 we present the results from a variety of specifications in which we estimate the effect of the boycott on various dependent variables that are related to retailers' marketing activities.

Each row in Table 6 represents a regression, differing only with respect to the dependent variable. The specifications are same as the difference-in-difference model in the previous analysis:

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

where  $Y$  represents the various dependent variables, and the remaining variables are as previously defined. The number of observations varies across regressions because particular dependent variables are missing for some observations (differing across specifications). As before, we use the data for the two months the boycott variable is switched on, combined with the same period of time one year before.

As reported in row (1) of Table 6, we find that the promotion price for French wine increases during the boycott by about 2.5%. Hence, retailers did not seek to mitigate the boycott by more aggressive discounting. In row (2) we report that the regular price also increased during the boycott, by slightly less than 1%. In rows (3) and (4) we show that the difference between the regular price and promotional price did not significantly change, in either absolute or proportional terms.<sup>13</sup> As we noted above, increasing prices may be explained by demand becoming more inelastic during the boycott.

In row (5) of Table 6 we report that the boycott caused a decrease in the fraction of French wine sold on promotion. This may be a demand response to higher promotional prices, although the promotional price relative to the regular price did not significantly change which suggests the

---

<sup>13</sup>The number of observations is less than row (1) because there are instances in which we observe the promotional price but not the regular price. A price for each category is only observed if there are strictly positive sales in that category.

relative demand may not change. If not a demand response to prices, an alternative explanation is that retailers reduced their marketing activities of French wine: less shelf space for promoted wines, less promotional advertising of French wine, and so forth. Since total sales of French wine fell during the boycott, and the fraction sold on promotion also fell, it follows that the quantity of French wine sold on promotion also decreases, as verified in row (6) of the table.

None of the evidence discussed in this subsection provides an ideal indication of whether retailers changed their behavior because of the boycott. But the available evidence suggests that retailers did not attempt to reverse the effects of the boycott by enhancing promotional activities of French wine. What is less clear is whether the retailers may have exacerbated the effects of the boycott by reducing their normal levels of promotional activity, or whether reduced promotional sales are a consequence of consumers' boycott participation. However, even if retailers did reduce their promotion of French wine, the magnitudes of the effects we describe in this section seem small in comparison to the overall reduction in demand we find in the other parts of our study. In other words, the evidence suggests that changes in retailers' behavior was probably not a major driver of the reduction in French wine sales.

### 3.4 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 5.0% and 15.3%. 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 exploits data to determine the start and end dates of the boycott.

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:

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

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 ( $\alpha_{0k}$ ), 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.<sup>14</sup>

There are two key parameters of interest. Firstly,  $\delta$  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  $0 < \delta < 1$ .<sup>15</sup> The closer that  $\delta$  is to one, the longer the boycott lasts. Secondly,  $\theta$  measures the contemporaneous response of consumers to the current call for a boycott (or more correctly:  $\theta$  measures consumers’ responsiveness to the depreciated stock of calls for a boycott). The more consumers that participate in the boycott, the larger the absolute value of  $\theta$ . With estimates of  $\delta$  and  $\theta$  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 7, including estimates for differing sub-samples and variations on the above specification. The top row of Table 7 is based on the full sample (4,160 observations). We obtain very precise estimates of both  $\delta$  and  $\theta$ :  $\hat{\delta} = 0.86$  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 2 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.

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 Table 7. 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

---

<sup>14</sup>The specific holidays are Valentine’s Day, Thanksgiving, Christmas, and New Years.

<sup>15</sup>This is not a constraint imposed for estimation.

$n_t = 0$  in all periods, and compute predicted quantities based on the estimated parameters. The time-series of the counterfactual is shown in Figure 2, 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*.<sup>16</sup> 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 would have been 13.3% higher over this six month period if there was no boycott. 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 7 we report the estimated duration to be 5.7 months, for the base specification.<sup>17</sup>

Figure 3 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 7 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 7, we find a larger maximum 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

---

<sup>16</sup>Factual sales are based on our estimated model, rather than the raw sales data.

<sup>17</sup>Note the dataset extends about 9 months after the start of the boycott.

sample to only European wines. In the third row of Table 7 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 7, the estimates for  $\theta$  and  $\delta$  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, French wine sales would have been approximately 13% higher during these six months if the boycott had not occurred, and at the peak of the boycott weekly sales would have been 27% higher if there was no boycott.

## 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.<sup>18</sup> The second characteristic we consider is whether buyers of cheap or expensive French wine were 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. In fact no city is 100% Republican or 100% Democrat. Each city is closer to 50-50 than these extremes, as shown in Table 2. Consequently, any difference in behavior between Republicans and Democrats needs to be quite dramatic if our analysis is to find significant differences.

---

<sup>18</sup>By 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.

We start by analyzing the variation in the boycott effect by political preferences. Figure 4 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. We re-estimated all the results in this study with the data for Boston excluded, and found no qualitative differences in any of the findings (recall that the Boston data captures a small share of the market).

For the remaining three cities in Figure 4, 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 city. The results are presented in Table 8. All coefficient estimates are significantly different from zero with 99% confidence.

We focus on the implied measures of the boycott magnitude and ignore the findings 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  $\theta$  and  $\delta$  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 5 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 9. All but one of the estimated coefficients are significantly different from zero with 99% confidence.<sup>19</sup> As Figure 5 indicated, the boycott of expensive French wines was intense but short-lived. We compute a maximum weekly effect of 52%, but a 6-month effect of only 12.4%.<sup>20</sup> 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 effect of 73.8% and a 6-month effect 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 a small (and statistically insignificant) 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).<sup>21</sup> The

---

<sup>19</sup>The estimate for  $\hat{\theta}$  in specification (3) is insignificantly different from zero at the 1% level.

<sup>20</sup>As throughout the paper, these magnitudes represent the predicted proportional increase in sales if there was no boycott, using the level of sales with the boycott as the base.

<sup>21</sup>In Figure 5, high price French wine is the only category to display a dramatic spike in sales in the holiday period.

reasoning is that buyers of the wine do 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.

## 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, there tend to be reasonably close substitutes for French wine. But this 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.<sup>22</sup>

---

<sup>22</sup>Prior research suggests the news media may be an important factor: Della Vigna and Kaplan (2007) shows that news coverage, and the Fox News Channel in particular, impacts on voting behavior.

To test the hypothesis 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 10 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  $\theta$  captures the instantaneous impact and the estimate for  $\delta$  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 10, it does not appear that Bill O’Reilly stimulated participation in the boycott.<sup>23</sup> Indeed, the point estimate of the instantaneous impact

---

<sup>23</sup>Recall, 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

of Bill O'Reilly is for an increase in sales of French wine (although it is insignificantly different from zero at the 95% confidence level). 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 10 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. These results support the view that news coverage is a proxy of the call for a boycott, as we assumed in the prior sections of this study.

## 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 the period of time during which there is a call for the boycott, 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 27% higher if there was no boycott. Also, over the six month period that we estimate the boycott lasted, 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

---

than 0.5.

French wine revenue of \$7.4 million, for Boston, Houston, Los Angeles and San Diego.<sup>24</sup> 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.<sup>25</sup> 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.

The boycott we have studied here can be described as a surrogate boycott, in which the French wine industry serves as a stand-in for the French government. There are many other kinds of boycotts, such as the common case of boycotts that target a particular company. It is conceivable that other boycotts would have more or less consumer participation than what we have found here. For instance, some consumers may be unwilling to participate in this boycott because they view the French wine industry as an arbitrary sub-group of the French population. Or consumers may be more or less responsive to other issues, such as animal rights or nuclear testing.

Our findings show that there can be significant consumer participation in a voluntary boycott, despite the free-riding problem economic theory suggests would inhibit participation. 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. We have not analyzed the potential impact on French government policy, although casual empiricism suggests there was no significant effect. This would be consistent with the findings of Tyran and Engelmann (2005) that show evidence of sellers' non-responsiveness to boycott activity in a laboratory experiment.

---

<sup>24</sup>Actually, 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.

<sup>25</sup>The 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.

The use of micro-level sales data reveals important variation in the degree of boycott activity: participation varies over time, across regions and across price-segments. The implication for managers is that a targeted response, such as targeted price reductions or advertising, can be an effective way of mitigating the harm to profits. Our analysis also highlights the usefulness of sales data for uncovering these patterns, rather than managers relying on surveys which can be an expensive alternative.

After our paper was written but prior to publication, there have been two other papers studying the French wine boycott. Ashenfelter, Cicarella and Shatz (2007) utilize sales data and estimate similar difference-in-difference models, but do not consider the nonlinear model we have emphasized. Their data differs from ours in a couple of key dimensions: their data is aggregate to the monthly and national level (ours is weekly by city), and their data ends in May 2003 which is before the end of the boycott according to our analysis (our data ends in November 2003). Contrary to our findings, they conclude there is no evidence of boycott participation. In a second paper, Bentzen and Smith (2007) find almost identical effects to our findings, based upon detailed but aggregate European export data.

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. Although, as a referee noted, the French government appears to have chosen a policy in spite of the actual damage to trade.

## Bibliography

- Adams Wine Handbook*, 2003, Adams Beverage Group, Norwalk, CT.
- Ashenfelter, O., S. Ciccarella and H. Shatz (2007): “French Wine and the U.S. Boycott of 2003: Does Politics Really Affect Commerce?” *Journal of Wine Economics*, 2(1), 55–74.
- 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.
- Bentzen, J. and V. Smith (2007): “The Military Action in Iraq 2003: Did U.S. Consumer Boycott of French Wines have any Economic Effects?” *Journal of Wine Economics*, 2(1), 75–83
- 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.
- Della Vigna, S. and E. Kaplan (2007): “The Fox News Effect: Media Bias and Voting,” *Quarterly Journal of Economics*, 122, 1187–234.
- 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.
- 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.

Tyran, J.R. and D. Engelmann (2005): “To Buy or Not to buy? An Experimental Study of Consumer Boycotts in Retail Markets,” *Economica*, 72, 1–16.

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
<b>TOTAL</b>	<b>1,099,858,923</b>		<b>130,627,884</b>		<b>8.42</b>	<b>14,175</b>

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	2 months	No	No	226,800	.86
(2)	$\ln(\text{Quantity})$	-.1213	0.0173	2 months, sales>0	No	No	107,821	.89
(3)	$\ln(\text{Quantity})$	-.1120	0.0173	2 months, sales>0	Yes	No	107,821	.90
(4)	$\ln(\text{Price})$	0.0113	0.0023	2 months, sales>0	No	No	107,821	.98
(5)	$\ln(\text{Quantity})$	-.1613	0.0207	2 months, sales>0, Euro-only	No	No	21,644	.88
(6)	$\ln(\text{Quantity})$	-.1665	0.0207	2 months, sales>0, Euro-only	Yes	No	21,644	.88
(7)	$\ln(\text{Price})$	-.0066	0.0029	2 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: Boycott effect for regions other than France (placebo regressions)

	Coeff.	Std. error	Coeff.	Std. error	Coeff.	Std. error	Coeff.	Std. error
<i>Boycott × France</i>	-.1166	0.0174	-.2942	0.0283	-.1154	0.0174	-.2940	0.0283
<i>Boycott × Italy</i>	0.0442	0.0150	-.0064	0.0221	0.0455	0.0150	-.0063	0.0221
<i>Boycott × Spain</i>	0.0175	0.0271	-.0130	0.0479	0.0187	0.0271	-.0129	0.0479
<i>Boycott × Germany</i>					0.1003	0.0443	0.0317	0.0346
Weights	No		Yes		No		Yes	
Observations	107,821		107,821		107,821		107,821	
$R^2$	0.89		0.93		0.89		0.93	

All regressions are based on the sample of wines for which we observe strictly positive sales in every wine-city pair. We also include wine-city fixed effects and week effects. Weights are given by the sales for city-wine pair.

Table 6: Difference-in-difference analysis of boycott effects on retailers' behavior

Dependent variable	Boycott coefficient	Standard error	Obs	$R^2$
(1) $\ln(\textit{Promotional Price})$	0.0250	0.0084	28,157	0.98
(2) $\ln(\textit{Regular Price})$	0.0094	0.0023	102,390	0.98
(3) $\textit{Regular} - \textit{Promotional Price}$	-0.2672	0.1530	22,726	0.51
(4) $(\textit{Regular Price} - \textit{Promotional Price})/(\textit{Regular Price})$	-0.0063	0.0089	22,726	0.45
(5) $\textit{Fraction Sold on Promotion}$	-0.0248	0.0054	106,165	0.48
(6) $\ln(\textit{Quantity Sold on Promotion})$	-0.1227	0.0143	106,165	0.91

All regressions include wine-city fixed effects and week fixed effects. Robust standard errors are reported. The number of observation varies across specifications because the promotional price variable is observed only if there are non-zero promotional sales.

Table 7: 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). Robust standard errors are reported. The 6-month effect is based on the period 2/10/03 to 8/17/03.

In the columns titled “max weekly effect” and “6-month effect” we report the predicted increase in sales of French wine if there was no boycott (i.e. the base is the level of sales with the boycott).

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.

Table 8: 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%	$\geq 8.5$
(2) Houston	-0.0150	0.0029	0.9455	0.0225	3,120	.99	22.3%	15.5%	$\geq 8.5$
(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

All regressions include holiday dummies for France, origin-specific time trends, and week fixed-effects. Robust standard errors are reported. The 6-month effect is based on the period 2/10/03 to 8/17/03.

In the columns titled “max weekly effect” and “6-month effect” we report the predicted increase in sales of French wine if there was no boycott (i.e. the base is the level of sales with the boycott).

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.

Table 9: 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	-0.0551	0.0072	0.8639	0.0281	3,328	.99	73.8%	34.4%	7.8
(2) \$7.99 – \$11.93	-0.0216	0.0011	0.9727	0.0084	2,912	.99	41.5%	30.5%	$\geq 8.5$
(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	-0.1162	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. Robust standard errors are reported. The 6-month effect is based on the period 2/10/03 to 8/17/03.

In the columns titled “max weekly effect” and “6-month effect” we report the predicted increase in sales of French wine if there was no boycott (i.e. the base is the level of sales with the boycott).

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.

Table 10: Effect of media on French wine boycott

Media source		Coefficient estimate	Standard error	Coefficient estimate	Standard error	Coefficient estimate	Standard error	Coefficient estimate	Standard error
Newspaper articles	$\hat{\theta}$	-0.238	0.0053	-0.206	0.0040				
	$\hat{\delta}$	0.8595	0.0468	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 French wine market share and boycott news articles




Figure 2: Actual, predicted and counterfactual French wine sales



Figure 3: Estimated percent increase in French wine sales if there was no boycott




Figure 4: French wine market share by city and boycott news articles





Figure 5: 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.