VOTING AFTER THE BOMBINGS: A NATURAL EXPERIMENT ON THE EFFECT OF TERRORIST ATTACKS ON DEMOCRATIC ELECTIONS

José G. Montalvo*

Abstract—Can terrorist attacks be timed to change the outcome of democratic elections? In this paper, we analyze the electoral impact of the terrorist attacks of March 11, 2004, in Madrid. Studies using individual level postelectoral survey data reach contradictory conclusions. We propose an alternative approach. Since the bombings took place only three days before the 2004 congressional election, we can find a control group of individuals who cast their vote before the terrorist attacks. The results indicate that the attacks had an important electoral impact, rejecting the hypothesis that the identity of the winner was unaffected by the terrorist attacks.

I. Introduction

ECONOMISTS have recently focused some efforts on analyzing the effects of terrorism and terrorist attacks in light of the September 11, 2001 attack on the United States. Several papers have shown that the timing of terrorist attacks is not random but corresponds to strategic objectives and political conditions. One much debated effect of terrorist attacks is their potential impact on the results of democratic elections if both events are close in time. During the 2004 U.S. presidential election, there was a broadly shared fear that terrorists would try to affect the election outcome using a terrorist attack.[1]

In this paper, we analyze the electoral impact of terrorist attacks using the March 11, 2004 bombings in Madrid, which happened three days before citizens voted. The popular press has argued that the objective of the attacks was to change the Spanish government since the incumbent party (conservative) had supported the war in Iraq, while the opposing (socialist) party was opposed. The conservative party lost the election, although the polls were quite favorable in the weeks before the election. Is it possible to find empirical evidence to support the claim that the terrorist

attacks affected the choice of the voters in the 2004 Spanish congressional election?

The first paper on this topic, by Lago and Montero (2005), argues, using data from postelectoral survey data, that the outcome of the Spanish congressional election of 2004 would have been the same without the terrorist attack ("Even if there would have been no terrorist attack, the most likely result is that the socialist party would have won the election"). By contrast Bali (2007), also using individual postelectoral survey data, concluded that the terrorist attacks "went a long way to explain the outcome of the election and, quite likely, decided it."

In general, it is difficult to evaluate properly the electoral effect of an episode like a terrorist attack using postelectoral surveys. Many problems can beset the results of this type of evaluation, for example, overrepresentation of the voters for the winning party or cognitive dissonance when asking retrospective questions.

In this paper we propose an alternative approach to evaluate the effect of the terrorist attacks on the electoral outcome; we use not postelectoral survey data but actual votes. Some special circumstances make the Madrid attacks suitable for a natural experiment. In particular, the attacks took place only three days before the election. The proximity of the bombings to voting divides the voters into two groups: one group knew about the terrorist attacks before they voted, and another group of voters did not know about them. The voters on the day of the election, March 14, knew about the terrorist attacks. But since the attacks took place so close to election day, it is possible to find out, and count, the actual votes of many voters who did not know about the terrorist attack when they voted: Spanish nationals abroad. We can use this special circumstance to interpret the data as a natural experiment and try to identify the effect of the attack on the results of the election using a difference-in-differences approach.

The consequences of the change in the outcome of an election may be very important. In the Spanish case, the winner of the election, the socialist party, decided to withdraw Spanish troops from Iraq immediately and undertook many reforms with important economic effects: changes in income and corporate tax codes, the regularization of immigrant workers, and changes in the structure of public expenditure, for example. These changes were not part of the conservative party's program.

Some political commentators, as well as part of the international press, argued that many of the votes lost by the conservative party after the attacks were due to the poor management of the crisis by the conservative government. Our identification strategy, based on actual votes rather than opinion polls, cannot separate these two effects (the terror-

[1] In early July 2004 the U.S. homeland security secretary, Tom Ridge, declared that al Qaeda intended to launch a "large-scale attack" inside the United States timed to the election to "disrupt our democratic process." Homeland Security even asked the Justice Department to consider what steps would need to be taken to postpone the presidential election in the event of a terrorist attack. On July 17, 2004, the *New York Times* ran an editorial against that proposal ("A Bad Idea, Rejected"). It reported on October 24 that "intelligence and counterterrorism officials say that they still fear an undetected plot, before the November 2 election. Others say they suspect a plan might have existed and been disrupted or postponed" (Johnston & Van Natla, 2007).

ist attack and the management of the information by the government after the attack). However, in the absence of terrorist attacks, how the government handled the information would have been irrelevant. Therefore, when we refer to the impact of the terrorist attacks, we consider a composite event that includes the bombings and how the conservative government handled information about the perpetrators of the attack.[2]

The paper is organized as follows. Section II describes the months previous to the election and the episodes that led up to election day. This section also discusses previous findings on the electoral impact of the terrorist attacks based on postelectoral surveys. Section III sets out our methodological strategy and discusses the main results of the estimation using a difference-in-differences estimator. Section IV presents a test of the robustness of the results using a synthetic control method for comparative case studies. Section V discusses some issues related to interpreting the results. Section VI contains the conclusions.

II. The Electoral Impact of the Madrid Terrorist Attacks

There has been a long and intense controversy about the electoral effect of the March 11 bombings in Madrid. In this section, we describe the events leading to the Spanish elections of March 14, 2004, and analyze several studies that have used postelectoral surveys to evaluate the impact of the attacks.

The political context is important to understanding the impact of the bombing. The relevant events started well before the attacks. In March 2003, the president of the Spanish government, José Maria Aznar, from the conservative party, joined the United States in the war against terrorism in Iraq. In October that year, Osama Bin Laden included Spain among the terrorist objectives of al Qaeda. The message he sent to Al Jazeera included the following sentence: "We have the right to retaliate whenever and wherever we believe it is adequate against the countries involved in the international coalition against Iraq, specially the United Kingdom, Spain, Australia, Poland, Italy and Japan." On March 11, 2004, thirty months after the 9/11 attack, Islamic terrorists deposited nine backpacks full of explosive in several trains in Madrid. The explosions killed 191 people and wounded 1,500. On March 14, the socialist party won the Spanish congressional election.

After the terrorist attack, the conservative government pointed to the ETA, the Basque terrorist group that seeks the independence of the Basque country, as the author of the attack. However, very soon there were many hints in the direction of radical Islamic groups, yet the government continued to argue that ETA was most likely responsible. By

[2] Bali (2007) argues, using the answers of a postelectoral survey, that the impact of government's handling of information on the attacks had a minor effect on the final result. Section V discusses this issue at length.

TABLE 1.—MONTHLY AVERAGE PERCENTAGES OF VOTES FOR DIFFERENT POLLS

	% Conservative	% Socialist	Difference
January	42.92	36.32	6.60
February	42.90	36.35	6.55
Before March 7	42.36	37.60	4.76
Election March 14	37.71	42.59	−4.88

January: Surveys of Sigma Dos, Noxa, Opina, Vox Publica, and Gallup. February: Surveys of Sigma Dos, Celeste Tel, Noxa, Metra Seis, Gallup, Citigate Sanchis, Demoscopia, CIS, and Ipsos-Eco Consulting. March: Surveys of Opina, Sigma Dos, Noxa, Vox Publica, Citigate Sanchis, Celeste Tel, Sigma Dos, and Demoscopia.

FIGURE 1.—EVOLUTION OF VOTING INTENTION
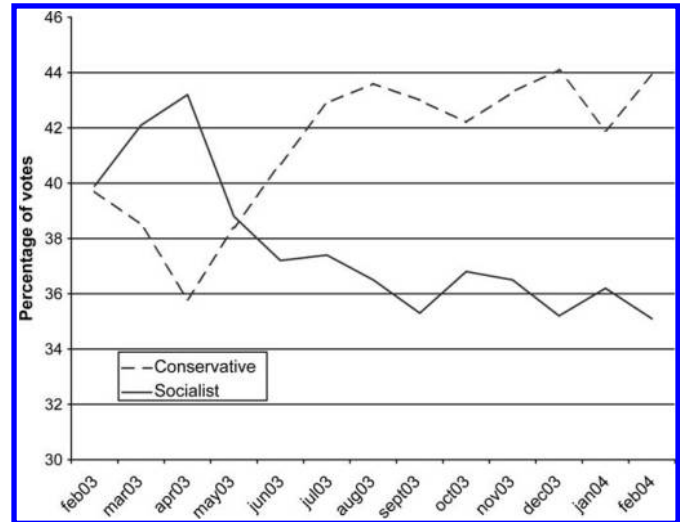
Gallup.

the afternoon of March 13, it was already quite clear that the attack had been executed by an Islamic terrorist group.

Was the terrorist attack the cause of the defeat of the conservative party? Polls taken during the legal period (until one week before election day) indicated that the conservative party had a clear advantage. Table 1 shows that the average percentage of vote for the conservatives was around 42% to 43% in the months prior to the election. Bali (2007) reports an average margin of advantage for the conservative party of 5.8 percentage points using seventeen polls in February and March. Therefore, the polls gave an important advantage to the conservative party. In fact as late as March 7, the only question was whether the conservative party would be able to get an absolute majority.

Table 1 shows averages of percentages from many different sources, some of them with a clear political preference for one of the largest parties. Figure 1 comes from a unique source, the Gallup poll, but shows a similar picture of citizens' voting intention. The data do not show any trend in the vote for the socialist party in the months prior to the election (March 2004). If anything it seems, from figure 1, that the conservative party was gaining some momentum during February. In fact, the last Gallup poll showed that 44% of the voters were planning to vote for conservative party candidates.

So, did the terrorist attack have a significant impact on the outcome of the election? Some authors say no. Lago

and Montero (2005), who used the Demoscopia postelectoral survey, argued that the attacks had no significant impact on the election results. They estimate a logit model for the probability of voting for the conservative party versus the socialist party using the individual microdata of Demoscopia. To construct these counterfactuals, they consider that nobody thought the attack was the result of Spain's foreign policy and nobody was influenced by how the government handled information on the attacks. Under these conditions, the conservative party would have increased only 1.2 percentage points over the actual percentage of vote obtained in the election. Therefore, Lago and Montero (2005) conclude that the terrorist attacks had only a limited effect.[3]

Bali (2007) reached the opposite conclusion. She also estimates a logit model but using the data of the CIS (the official Spanish Center for Sociological Research) postelectoral survey. The idea of the counterfactual is similar to the one in Lago and Montero (2005). Therefore, she also fixes the value of certain variables to 0 in her logit equation to estimate voting results in the absence of the bombings. Initially, considering that those encouraged to switch their vote did not switch and that those encouraged to turn out did not turn out, Bali predicts an increase of the share of the vote of the conservative party of 5.5 percentage points. A second exercise, using a preelectoral sample, leads to a counterfactual increase in the share of the conservative party of 8 percentage points. Summarizing, Bali (2007) concludes that "the attack had a considerable impact on vote decisions."

Estimating the counterfactual (no terrorist attacks) percentage of votes for the conservative party using alternative postelectoral surveys leads to very different results. Lago and Montero (2005) conclude that the terrorist attack had practically no influence on the results of the elections, while Bali (2007) reaches quite the opposite conclusion. However, the use of postelectoral surveys for this type of counterfactual analysis is problematic. First, such surveys tend to overrepresent the winner. For instance, the Demoscopia sample is clearly biased since it grossly underrepresents the voters for the conservative party. In this poll, there are almost two voters for the socialist party for each voter for the conservative party when, in the actual election, the ratio was close to half.[4] In addition, when a postelectoral survey asks for the reasons that voters changed their choice after this shocking event, the answers may be conditioned by problems associated with retrospective questions and cognitive dissonance. In this case, the problems are far more important since the terrorist attack could have a critical impact on the recollections and perceptions of what voters thought before the attack. Bali (2007) warns that some

respondents who claimed that the terrorist attack had no influence on their decision may actually have been affected despite their reports and may have acted differently had the terrorist attack not taken place. For all these reasons, using a postelectoral survey to analyze counterfactuals related with the electoral impact of the March 11, 2004 terrorist attacks could be problematic.

III. Analyzing the Counterfactual: An Alternative Identification Strategy

Is there an alternative to the use of postelectoral survey data for analyzing the counterfactual? Data on actual votes can solve some of the problems associated with the analysis of the counterfactual using individual survey data. For instance, Wand et al. (2001) analyze the electoral impact of the butterfly ballot used in Palm Beach County, Florida, in the 2000 presidential election. Palm Beach County is a traditionally Democratic and liberal county. However, Pat Buchanan, a conservative candidate, had an unexpected and impressive number of votes in the 2000 presidential election. Many analysts claimed that the confusing format of the butterfly ballot used in Palm Beach County misled many Democratic voters to vote inadvertently for Buchanan. Wand et al. notice that since the butterfly format was not used for absentee ballots, the election generates a natural experiment: the election day voters used the butterfly ballot, while the absentee voters did not. Therefore, if the butterfly ballot was responsible for the very high percentage of votes for Buchanan, it should be that his support came disproportionately from election day voters. Wand et al. show that Buchanan's proportion of vote on the election day ballot was four times larger than his proportion on the absentee ballot. However, this difference estimator has limitations: the mechanism that allocates voters to the election day or the absentee pool is not a random assignment. Wand et al. assume that even if absentee voters are not representative of election day voters, their political preferences are similar across counties in Florida. Since there was not a significant difference in the proportion of votes for Buchanan between election day voters and absentee voters in the rest of the counties of Florida, Wand et al. conclude that the butterfly ballot was pivotal to the 2000 presidential election. Al Gore would have won the majority of certified votes had Palm Beach County not used a confusing ballot.

A. The Identification Strategy

The identification strategy proposed in this paper also relies on an interpretation of the terrorist attacks as a natural experiment. The case of the Madrid bombings is a special situation because of their timing and the voting deadlines for Spanish nationals living abroad. This timing generates a natural experiment that relies only on observed voting behavior. The voters who cast their vote physically on March 14 obviously knew about the terrorist attack. However, "absent

---

[3] Torcal and Rico (2004) also support this claim, although with a weaker set of results than Lago and Montero (2005).

[4] The degree of overreporting for the winner in the surveys that Bali (2007) used is smaller than the one on the Demoscopia poll that Lago and Montero (2005) used.

[voters] residents abroad" could vote before the election date.[5] In many countries, the deadline for citizens who are voting abroad is several days before the official day of the election.

Notice that we are not referring to voters by mail in general. The vote by residents in Spain not in person cannot be differentiated from the vote in person because it is cast in the same urn.[6] The absent (voters) residents abroad, who live in a foreign country and are registered with a Spanish consulate, vote by mail. The electoral rules for absentee voting in the 2004 elections allowed them to vote before the day of the bombing. In particular, they could start voting on March 2 and had two possible options: casting their vote in person before March 7 at a Spanish consulate or by certified mail.[7] We can use two groups, Spanish residents and Spanish nationals abroad, to evaluate the effect of the terrorist attack on the results of the election. Using the terminology of causal inference, we have two groups: the treated group is the set of voters resident in Spain, who were subject to the treatment (information about the terrorist attack) before they voted, and the control group, which consists of the absentee voting by Spanish nationals abroad, who did not know about the terrorist attack before they cast their vote.

It is reasonable to suspect that Spanish residents abroad are not representative of the total population of Spanish voters in congressional elections. However, and in contrast with Wand et al. (2001), we use the time-series variation on both groups to identify the effect of the terrorist bombing. Wand et al. assume that even if absentee voters are not representative of all voters, the political preferences of absentee voters were similar in all counties in Florida. Instead of using this type of cross-sectional assumption for the identification of the effects, we rely on the usual assumption of parallel trend that justifies the difference-in-differences estimator.

For both groups (election day voters and Spanish residents abroad) we have the results of several congressional elections. For identification, we need the assumption that in the absence of treatment, the average outcome for the treated and the untreated would have followed parallel trends. Figure 2 shows the evolution of the average ratio of voters of the conservative party over the socialist party in the provinces of the treatment and control groups. Before 2004,

FIGURE 2.—AVERAGE RATIO ACROSS PROVINCES OF VOTES OF CONSERVATIVES OVER SOCIALISTS
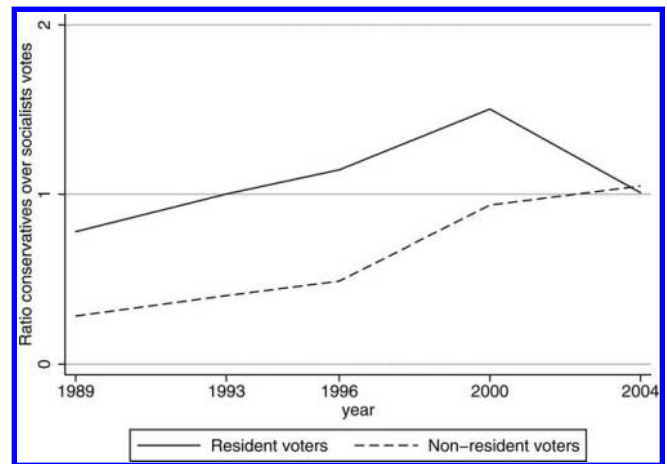
the lines are basically parallel; in 2004, they converge. As expected, election day voters are quite different from the Spanish residents abroad in their political preferences. Nonetheless, the difference-in-differences setup should take care of this difference. Figure 2 indicates that the common trends assumption is not at odds with the data.

B. Data and Basic Results

Our data contain 52 observations for each time period for the treated group (Spanish residents) and the control group (Spanish citizens in a foreign country). These observations correspond to the results of all the congressional elections between 1989 and 2004 in each of Spain's 52 provinces.[8] All data are from the database of the Spanish Ministry of Interior (Ministerio del Interior). The electoral roll is divided into two parts: the regular electoral roll and the electoral roll of "absent (voters) residents abroad," or voters who have registered with a Spanish consulate because they live abroad (CERA, or "Electoral Roll of Absent Residents Abroad").

Consider the specification

$$Y_{igt} = \mu + \lambda_t + \beta G_g + \delta(DG)_{gt} + v_{gt} + u_{igt}, \quad (1)$$

---

[5] In Spain these voters are counted in the CERA (Electoral Roll of Absent Residents Abroad). These persons are basically Spanish nationals with the right to vote who are living in foreign countries and are registered at Spanish consulates around the world.

[6] Postal voters who live in Spain but are not going to be in their electoral district on election day are not counted in CERA. Strictly speaking, these voters are also "absent voters," but they are not included in the CERA because their official residence is a Spanish address. In principle, the resident postal voters contaminate the treated group (election day voters) since they did not know about the terrorist attack when they voted, but they are counted together with the election day voters who knew about it. However, resident postal voters represent less than 2% of the non-CERA votes, so the impact of this small contamination on the estimates will be very limited.

[7] A caveat we discuss later is the possibility of other important news (apart from the terrorist attack) that changed electors' voting choice between March 2 and election day.

[8] We consider only congressional elections since voters' choices in other type of elections (local, regional, or European) are significantly different. The first election considered is that of 1989. For the elections before 1986, there is no information on votes of the Spanish residents abroad by electoral districts. The election of 1989 is also the first election in which the Popular party participates under this name, and with a new political vision that tries to break with the image of the old coalition, Alianza Popular, as the heritage of the Franco regime. In addition, the demographics of refugees from the Civil War living abroad in the initial years of the Spanish democracy, and their political bias, is likely to have generated differential effects in the voting of Spanish abroad with respect to the residents as a reaction to rapid changes in the political arena during the initial years of the Spanish democratic experience (for example, dissolution of the centrist party UCD, political stigma of Alianza Popular, etc.).

TABLE 2.—BASIC ESTIMATION

	Entire Sample	Placebo (Attacks date = 2000)
Year 2004	0.80	
	[0.04] ***	
	(0.03) ***	
	{0.06} ***	
Year 2000	0.68	0.69
	[0.04] ***	[0.05] ***
	(0.03) ***	(0.03) ***
	{0.05} ***	{0.07} ***
Year 1996	0.28	0.28
	[0.02] ***	[0.02] ***
	(0.04) ***	(0.04) ***
	{0.05} ***	{0.06} **
Year 1993	0.17	0.17
	[0.01] ***	[0.01] ***
	(0.03) ***	(0.03) ***
	{0.05} **	{0.06} *
Resident	0.57	0.58
	[0.04] ***	[0.04] ***
	(0.02) ***	(0.03) ***
	{0.03} ***	{0.05} ***
Resident × Year 2004	−0.61	
	[0.04] ***	
	(0.02) ***	
	{0.08} ***	
Resident × Year 2000		−0.01
		[0.05]
		(0.03)
		{0.09}
Constant	0.24	0.24
	[0.03] ***	[0.03] ***
	(0.03) ***	(0.04) ***
	{0.04} ***	{0.04} ***
R^2	0.46	0.51
H_0: ratio $< 1 | \delta = 0$	$p = 0.00$	
N	520	416

In brackets, pooled OLS clustered (at the province level) robust standard error, and in parentheses, pooled OLS clustered (at the group/time level) robust standard error. In braces, Donald and Lang's (2007) estimator of the standard error. The coefficient estimated is the same for all three procedures. ***Significant at 1%. **Significant at 5%. *Significant at 10%.

where i ($i = 1, \ldots, I$) indexes provinces, g indexes groups (treatment or control), and t indexes time. The province is the highest level of disaggregation since the electoral districts for Spanish nationals abroad are the provinces. Y is the outcome variable (in our case, the ratio of the vote for the conservative over the socialist party), G is a dummy variable that denotes the group ($G = 0$ if the observation belongs to the control group and $G = 1$ if it belongs to the treatment group), λ is a time dummy variable, DG is a group/time period covariate (it takes the value 1 for the treatment group after the terrorist attacks), v_{gt} is un unobserved group/time effect and u_{igt} is an error term. We are interested in estimating the parameter δ. A sufficient condition for identifying the treatment effect is that the selection for treatment does not depend on the individual-transitory perturbation of the primitive specification.

Table 2 presents the estimation of specification (1). A common practice for estimating this specification is to ignore the unobserved group/time effect. This amounts to assuming that individual-level observations are independent. Column 1 reports the pooled OLS estimator of the coefficients. The first row after each parameter estimate includes the clustered (at the provincial level) robust stan-

dard error.[9] When this initial estimator of the standard deviation is used, the coefficient for the treatment group after the treatment is highly statistically significant. The potential problem with the group error terms derives from the fact that the unit of observation (province) is more detailed than the level of variation (treatment and control groups over time). The presence of common group error terms can generate estimated standard errors that are dramatically biased.

One possible solution is to use the pooled OLS with a clustered-robust standard error. In our case, we interpret each group/year combination as a cluster. We define those clusters as c ($c = 1, \ldots, C$) and the size of each cluster as M_c. Our data contain 10 of those clusters with 52 observations each. The size of all the clusters is the same, and therefore $M_c = M$. It is well known, by analogy with conventional panel asymptotics, that the clustered covariance matrix is valid for inference when C goes to infinity and the size of each group is fixed. Hansen (2007) extends this result and shows that the inference with the clustered-robust covariance matrix estimator, appropriately normalized, is consistent, even with arbitrary correlation among the unobserved errors within each group, when the number of clusters and the number of observations in each cluster get large. The second rows show the clustered-robust standard errors and the level of significance of the coefficients using those estimates for doing inference.[10] The estimated parameter δ, which measures the impact of the terrorist attacks, is also statistically significant using a clustered (at the group/time level) robust standard error.

However, the results in Hansen (2007) do not justify cluster robust inference using pooled OLS when the number of groups is small.[11] In our case, the pooled OLS is estimated with a small C (in fact, ten clusters) and a relatively large group size. In fact, the properties of the cluster-robust inference based on a small number of groups and a large group size when cluster effects are left in the error term are probably not very good. In brief, fully robust inference is justified for large C and small M and when M is large and C is also reasonably large, but is likely to produce poor inference if C is small, even if we do not leave the unobserved (clustered) effects in the error term.[12]

---

[9] Since there are no individual specific covariates, the pooled OLS estimator leads to the same results as the fixed-effects estimator with a robust standard error.

[10] The estimated coefficients are the same.

[11] Hansen (2007), theorem 4, shows that in general, the typical cluster-robust covariance is not consistent when the number of groups is fixed and the size of each group goes to infinity. Under additional restrictions, Hansen shows that the limiting distribution of the cluster-robust covariance estimator is proportional to the covariance matrix of the data. This is convenient because in the panel case, the normalization by $C/(C - 1)$ will lead to inference based on a t_{C-1} distribution.

[12] Despite this theoretical results, the simulations reported in Hansen (2007) for $C = 10$ and $M = 50$, which defines a setup very similar to ours, show that the cluster-robust variance estimator with critical values from a t_{C-1} distribution produces very small distortions. However, the results of this simulation exercise could be compromised by the specific parameters used in the simulation.

Therefore, we need to check if the results are robust to inferential methods appropriate for C small and M relatively large. This problem was initially studied by Moulton (1990). Donald and Lang (2007) considered the estimation in the context of a small C and large M. They argue that failing to take into account the group error structure will bias the standard errors; moreover, using a normal distribution to the corrected t-statistics overstates the significance of the estimated coefficients. Donald and Lang propose a solution to this problem. The estimated coefficient, if the size of all the groups is the same, is identical to the pooled OLS and the random effects estimator. The basic difference is the estimation of the standard error of those parameter estimates. In order to do inference, Donald and Lang assume that the conditional distributions of the group error term and the individual-specific terms are normally distributed. Under this assumption, the inference can be performed using a t_{C-K} distribution, although we can also obtain asymptotic normality if the individual-specific effects are not normally distributed provided that the number of observations in each group is large and the group error term is normally distributed conditional on the between-group varying variables. If C is small, then using that distribution with $C - K$ degrees of freedom is stricter than using the t_{MC-K} distribution derived from the standard inference for pooled OLS. The third row after each estimator of the coefficients shows the standard errors and the level of significance of the estimated coefficients using the inferential results derived from the approach of Donald and Lang. Using these standard errors on the estimated parameter of the effect of the bombings multiplies by 4 the estimated cluster-robust standard error from the pooled OLS estimator with clustered (at the group/time level) robust standard error. Nevertheless, the estimator of δ is still statistically significant using this more stringent criterion.

Therefore, the different estimators of the standard deviation of the estimator lead to the same conclusion: the estimated coefficient of the effect of the bombings is statistically significant.[13] We now return to the common trend hypothesis, the basic assumption for the identification of the effect. To examine further the likelihood of the common trend hypothesis, we run a placebo experiment. We estimate the basic specification on a placebo bombings taking place in the year 2000. For this estimation, we discard the year 2004 from the data. Column 2 shows the results of the estimation. No matter what calculation of the standard error we choose to make inferences (cluster-robust pooled OLS by country or by group/year or the Donald and Lang approach) the result is the same: the coefficient on the placebo bombings is not statistically significant. The similarity of the estimators of the parameters in columns 1 and 2 is also remarkable except, obviously, for δ.

Using the results of the estimations, we can test if the impact of the terrorist attack changed the identity of the winning party. Notice that we still understand this impact as the combined effect that the terrorist attack had on voters' choices because of several reasons that we cannot identify separately (the recollection of government's support for the Iraq intervention and the influence of the belief that the incumbent party had mishandled the postattack events by manipulating or hiding information). We want to test the null hypothesis of the victory of the socialist party, that is, a ratio smaller than 1, under the counterfactual situation of no terrorist attacks. The one-sided test of this hypothesis rejects the null at the 1% level of significance no matter what type of standard errors are used to perform the test.[14]

The results of elections are usually presented in the form of percentage of votes for a political party. We can calculate the percentage of the votes for the conservative party under the no-attack counterfactual using the predicted ratio and some accessory assumptions. We initially calculate the predicted weighted ratio of conservative over socialist votes by setting the variable DG to 0. Then we divide the predicted ratio by 1 plus the predicted ratio. This generates the counterfactual ratio of votes for the conservative party over the sum of votes of the conservative and the socialist parties. In order to get the counterfactual percentage of total votes for the conservative party, we need to multiply that predicted ratio by the sum of the percentage of total votes for the conservative party plus the socialist party. If we use the actual sum of the percentage of votes for both parties, then the conservative party would have obtained 42.7% of the votes, 5 percentage points above the actual outcome of the 2004 election. Instead, we could predict the sum of the percentage of votes for the two parties using a time trend, since the support for the two largest parties has been growing over time. When we use this second estimate, the conservative party would have obtained 44.4% of the votes, amounting to an increase of 6.7 percentage points over the actual result of the election.[15]

IV. A Synthetic Control Method

The standard approach to inference in the difference-in-differences setup assumes that the uncertainty is associated with the sampling error in estimating the means of each group/time combination. However, researchers have suggested different sources of uncertainty (see Bertrand, Duflo, & Mullainathan, 2004; Donald and Lang, 2007; Hansen, 2007). More recently, Abadie, Diamond, and Hainmueller (2010) argue that, in comparative case studies, there is an additional source of uncertainty: the capability of the control group to generate the counterfactual evolution of the

---

[13] The same conclusion is obtained if we run a weighted regression in which we weight each observation by the proportion of voters in each electoral district.

[14] This null hypothesis is also rejected at the 1% level of significance if we use the results of the weighted regression.

[15] These two scenarios are compatible with the results of the transformation approach in Montalvo (2006).

outcome that the affected units would have followed in the absence of the intervention. For this reason, it is possible that alternative control groups may reflect the counterfactual evolution of the outcome variable in the absence of intervention. For instance, Card (1990) uses a combination of cities in the United States (Atlanta, Los Angeles, Houston, and Tampa–St. Petersburg) to approximate Miami's labor market in the absence of the Mariel boatlift, a boatlift of 125,000 Cuban refugees who left Mariel Harbor in Cuba in 1980 and immigrated to the United States. However, Card's justification for using that particular combination of cities as the control group is not formal or based on a specific methodology: "These four cities were selected both because they had relatively large populations of blacks and Hispanics and because they exhibited a pattern of economic growth similar to that in Miami."

By contrast, Abadie and Gardeazabal (2003) provide a specific, data-driven procedure to choose the elements of the control group in their study of the economic impact of terrorism in the Basque country. They use a combination of two Spanish regions to mimic the economic evolution of the Basque country in the absence of terrorism. Abadie et al. (2010) analyze the theoretical underpinnings for the construction of synthetic control groups and for doing inference in this setup. The construction of a rigorously based synthetic control group provides a transparent approach to the composition of the comparison groups and enables researchers to decide on the design without access to postintervention outcomes. Abadie et al. construct a simple econometric model that can justify the synthetic control approach under more general conditions than is usual for linear panel data and difference-in-differences estimators.

The objective of the synthetic control group methodology is to choose a vector of weights so as to minimize some distance between the preintervention characteristics of the exposed unit and the characteristics of the (weighted) synthetic control group.[16] We assume that X_1 is a vector of preintervention variables for the treated region, and X_0 is a matrix of characteristics of the unexposed regions. The objective is to chose a vector of weights, W, such that we minimize some distance function, $\|X_1 - X_0 W\|$. In particular, Abadie et al. consider the following distance function,

$$\|X_1 - X_{0W}\|_V = \sqrt{(X_1 - X_0 W)' V (X_1 - X_0 W)},$$

where V is some symmetric positive semidefinite matrix. The optimal choice of V is derived after assigning weights to linear combinations of X_1 and X_0 to minimize the mean square error.

As a robustness test for the results presented in the previous section, we construct a synthetic control group and compare it with the treated group. In our case, we want to apply the synthetic control methodology to construct ade-

TABLE 3.—COMPARISON GROUPS

	Treated	Synthetic
Percentage of votes for the conservative party in 1989	25.8	25.6
Percentage of votes for the conservative party in 1993	34.8	35.0
Ratio of voters of the conservative party over the socialist party in 1996	1.03	1.03
Ratio of voters of the conservative party over the socialist party in 2000	1.30	1.30

quate comparison groups for the aggregate evolution of the ratio of votes of conservatives over socialists using the pool of nonexposed voters by provinces. Therefore, the setup is slightly different from that of Abadie et al. (2010), where the comparison involves one region exposed to the treatment and a pool of nonexposed regions. In our data, we want to mimic the aggregate evolution of the outcome variable among the treated voters using the pool of provinces containing the nontreated voters as the basis for the construction of the comparison group. We can use the results of the comparison of the treated group and the synthetic control group as a robustness test for the results derived from the difference-in-differences estimator. As we have already argued, the synthetic control method generalizes the assumptions in the usual difference-in-differences estimator.[17]

Table 3 compares the pretreatment characteristics of the ratio of conservative to socialist voters in the treatment group (resident voters) with the ratio in the synthetic control group (nonresident voters).[18] The characteristics used for the prediction are lags of the percentage of votes for the conservative party (1989 and 1993) and lags of the ratio (1996 and 2000).[19] The synthetic control group is similar to the treated group. However, it is quite different from the average of the 52 control provinces. In the previous section, we showed that the average ratio is very different in the nontreated group (Spanish nationals abroad). The percentage of vote for the conservative party in the average of the 52 control provinces (13.7% in 1989 and 23.5% in 1993) is also quite different from the share of votes for the conservative party in the treated group (election day voters) and the synthetic control group.

Figure 3 plots the trends in the actual ratio and the ratio for the synthetic group. Despite the small number of pretreatment periods and the use of a few characteristics, the

---

[16] The method resembles the procedures of matching.

[17] Athey and Imbens (2006) provide also a generalization of the difference-indifferences model for a nonlinear set up.

[18] For the calculations in this section, we have used the publicly available STATA routine written by Jens Hainmueller.

[19] A small number of control variables can be used for this exercise since the information on the characteristics of the voters abroad is quite limited. Despite this limitation, the adjustment before the year of the terrorist attacks is good. The combination of controls considered produces the smallest preintervention root mean square error. The results are very similar if we consider the percentage of votes of socialists in 1989 instead of the conservative vote in 1993. In fact, the root-mean-square error in this case is 0.05, which represents around 5% of the average value of the outcome variable.

FIGURE 3.—TRENDS IN THE RATIO OF CONSERVATIVE TO SOCIALIST VOTES: TREATED AGGREGATE GROUP VERSUS SYNTHETIC CONTROL GROUP
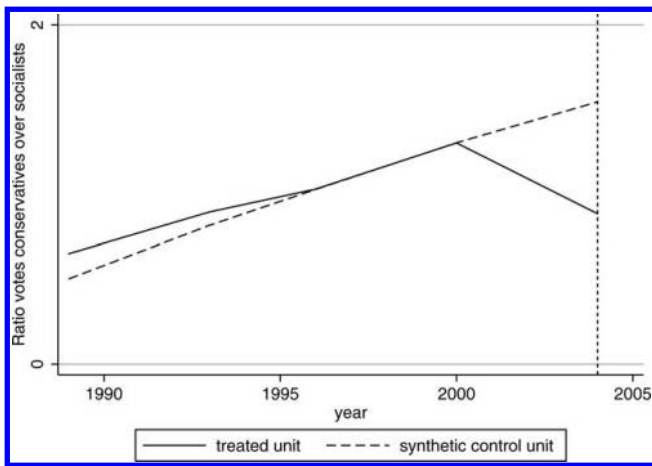
synthetic control group closely resembles the actual ratio. The root mean square prediction error is 0.08, which represents around 8% of the average value of the outcome variable. Most important, the effect of the terrorist attacks is similar to the estimation in section III. In the year of the terrorist attacks, the difference between the treatment and the synthetic control group is around -0.55 points. This result is similar to the estimate obtained using the difference-in-differences estimator (-0.61).

V. Discussion

The results of the DID estimation presented in section III show that the terrorist attacks had a statistically significant negative effect on the electoral support of the conservative party in the 2004 congressional election. This conclusion is robust to the use of the synthetic control method for comparative studies proposed by Abadie et al. (2010). The null hypothesis of a victory of the socialist party in the absence of the terrorist attack is rejected by the data. In addition, the estimated counterfactual percentage of votes for the conservative party ranges, depending on the sustained assumption on the joint percentage of votes for the two largest parties, between 42.7% and 44.4%. These findings are substantive since the actual percentage of votes for the conservative party was 37.4%. Note also that the socialist party won the election, and the presidency of the government, with a percentage of votes similar to our estimated counterfactual share of votes for the conservative party. The plausibility of the results of this article is supported by the results of the 2004 preelectoral polls. The lower bound of the range of counterfactual votes for the conservative party is similar to the average of the polls presented in table 1, while the upper bound corresponds closely to the latest Gallup poll presented in figure 1. Obviously this is just an indicator of the plausibility of the results, not a confirmation of the findings in any way. Preelectoral polls do not always predict actual electoral results accurately. In fact, one of the objectives of

the methodology proposed in this paper is to avoid the use of any type of surveys and polls in the evaluation of the impact of the attacks.

That Spanish nationals abroad could note as early as March 2, whereas the terrorist attack was on March 11, could also be a source of threats to the internal validity of the results. If there were negative news about the conservative party or positive news about the socialist party between March 2 and the terrorist attacks, then the treatment effect estimated previously could not be fully attributed to the terrorist attacks. We checked the best-selling newspapers of Spain during the period between March 2 and 11 searching for negative news about the conservative party (for example, corruption scandals, worsening of the prospects of victory in the congressional election).[20] A great deal of news is related to the two largest political parties (conservative and socialist). In analyzing more than 200 news items during that period, we were unable to find any clear bias for or against the conservative party. The fact that there was no important news during the period, that almost every newspaper has a clear bias for or against one of the two largest parties, and that the circulation of newspapers favoring each option is quite balanced can explain this absence of aggregated bias over the period and the set of chosen newspapers.[21]

Finally, there has been a long discussion on the likely reasons for the electoral impact of the terrorist attacks. The literature and the popular press emphasized basically two explanations: the discontent with the government of the conservative party for its support of military intervention in Iraq and disapproval of the government's handling of the information about the early stages of the investigation. Bali (2007) argues that "the concerns with foreign policy and government transparency exhaust the ways a respondent [in the postelectoral survey] was influenced by the attack." Bali concludes, using the OPINA postelectoral survey data, that those whose vote was influenced by the conservative party's policy on Iraq were more likely to vote for the socialist party, by 62 points, and the belief that the president of the conservative government mishandled the information about the authorship of the attacks translated into support for the socialists by 24 points. Therefore, it seems that the OPINA postelectoral survey indicates a predominant effect of the policy on Iraq over the manipulation, or hiding, of information.

Obviously our methodology, based only on actual votes, cannot provide any insight into the specific reasons behind the electoral impact of the terrorist attacks. There is a trade-off between the methodologies used to evaluate the impact

[20] We included all the newspapers with a net average daily circulation of over 100,000 copies (except sports newspapers). The list includes *ABC* (278,000), *El Correo* (124,999), *El Mundo* (314,000), *El País* (453,000), *El Periódico* (170,000), *La Razón* (140,000), and *La Vanguardia* (208,000). Source: OJD (Institutional Supervisor of circulation in written mass media). Net average daily circulation is shown in parentheses.
[21] We also checked the same for the 1996 and 2000 congressional elections without finding any significant event.

of the terrorist attacks. The analysis based on postelectoral surveys allows, if the appropriate questions are asked, the study of the reasons behind the change in voting behavior due to the attacks. However, the fact that surveys are based on retrospective questions asked of individuals limits the credibility of the results concerning the overall electoral impact of the bombings. The natural experiment approach proposed in this paper can identify more credibly the overall electoral impact of the attacks, but it is silent about the reasons behind the influence of the attacks on voters' choices.

VI. Conclusion

Recent papers (for instance, Pape, 2003) show that the timing of terrorist activities is not random but corresponds to strategic objectives and political conditions. Could terrorist attacks be timed to alter the result of democratic elections? This paper looks at the electoral effect of terrorist attacks in Madrid in 2004. One week before the election, opinion polls showed a clear advantage for the conservative party. After the terrorist attacks, the socialist party won the congressional election.

The impact of the bombings generated heated public debate. There was also controversy on the scientific front. Two papers, using individual data from several postelectoral surveys, have reached opposite conclusions. However, the use of postelectoral surveys could be problematic for analyzing the counterfactual question of what would have happened if the terrorist attacks had not taken place. This paper proposes an alternative approach to identifying the impact based on the interpretation of the attacks as a natural experiment, and a difference-in-differences estimator, to evaluate the electoral effect of the March 11, 2004 bombings in Madrid. The placebo experiment indicates that the parallel trend assumption is not at odds with the data. The calculations under the counterfactual of "no terrorist attacks" are in line with polls taken prior to the attacks. The findings are robust to the use of alternative estimators and methodologies (like the synthetic control methods). The results reject the null hypothesis that the socialist party would have won the election in the absence of the bombings.

REFERENCES

"A Bad Idea, Rejected," *New York Times*, July 17, 2004.
Abadie, A., A. Diamond, & J. Hainmueller, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effects of California's Tobacco Control Group," *Journal of the American Statistical Association* 105 (2010), 495–505.
Abadie, A., & J. Gardeazabal, "The Economic Cost of Conflict: A Case Study of the Basque Country," *American Economic Review* 93 (2003), 113–132.
Athey, S., & G. Imbens, "Identification and Inference in Nonlinear Difference-in-Difference Models," *Econometrica* 74 (2006), 431–497.
Bali, V., "Terror and Elections: Lessons from Spain," *Electoral Studies* 26 (2007), 669–687.
Bertrand, M., E. Duflo, & S. Mullainathan, "How Much Should We Trust Differences in Differences Estimates?" *Quarterly Journal of Economics* 119 (2004), 249–275.
Card, D., "The Impact of the Mariel Boatlift on the Miami Labor Market," *Industrial and Labor Relations Review* 44 (1990), 245–257.
Donald, S., & K. Lang, "Inference in Difference-in-Differences and Other Panel Data," this REVIEW, 89 (2007), 221–233.
Hansen, C. J. "Asymptotic Properties of a Robust Variance Matrix Estimator for Panel Data when T Is Large," *Journal of Econometrics* 141 (2007), 597–620.
Johnston, D., and D. Van Natla, "The 2004 Campaign: Security," *New York Times*, October 24, 2007.
Lago, I., and J. R. Montero, "The Mechanics of Electoral Change," *Claves de la Razón Práctica* 149 (2005), 36–44.
Montalvo, J. G., "Voting after the Bombing: Can Terrorist Attacks Change the Outcome of Democratic Elections?" Universitat Pompeu Fabra working paper 1000 (2006).
Moulton, B., "An Illustration of Pitfall in Estimating the Effects of Aggregate Variables on Micro Units," this REVIEW 72 (1990), 334–338.
Pape, R., "The Strategic Logic of Suicide Terrorism," *American Political Science Review* 97 (2003), 343–361.
Torcal, M., and G. Rico, "The Spanish General Election: In the Shadow of Al-Qaeda," *Southern European Society and Politics* 9 (2004), 107–121.
Wand, J., K. Shotts, J. Sekhon, W. Mebane, M. Herron, and H. Brady, "The Butterfly Did It: The Aberrant vote for Buchanan in Palm Beach County, Florida," *American Political Science Review* 95 (2001), 793–810.

This article has been cited by:

1. Martin Karlsson, Stefan Pichler. 2015. Demographic consequences of HIV. *Journal of Population Economics* **28**, 1097-1135. [CrossRef]

2. Simone Schüller. 2015. The 9/11 conservative shift. *Economics Letters* **135**, 80-84. [CrossRef]

3. Paolo Pinotti. 2015. The Economic Costs of Organised Crime: Evidence from Southern Italy. *The Economic Journal* **125**:10.1111/ecoj.2015.125.issue-586, F203-F232. [CrossRef]

4. Eric D. Gould, Esteban F. Klor. 2015. The Long-run Effect of 9/11: Terrorism, Backlash, and the Assimilation of Muslim Immigrants in the West. *The Economic Journal* n/a-n/a. [CrossRef]

5. Woo Chang Kang. 2015. Electoral cycles in pork barrel politics: Evidence from South Korea 1989–2008. *Electoral Studies* **38**, 46-58. [CrossRef]

6. Yusaku Horiuchi, Asher Mayerson. 2015. The Opportunity Cost of Conflict: Statistically Comparing Israel and Synthetic Israel. *Political Science Research and Methods* 1-10. [CrossRef]

7. Ho Fai Chan, Bruno S. Frey, Jana Gallus, Benno Torgler. 2014. Academic honors and performance. *Labour Economics* **31**, 188-204. [CrossRef]

8. Sivan Hirsch-Hoefler, Daphna Canetti, Carmit Rapaport, Stevan E. Hobfoll. 2014. Conflict will Harden your Heart: Exposure to Violence, Psychological Distress, and Peace Barriers in Israel and Palestine. *British Journal of Political Science* 1-15. [CrossRef]

9. ANNA GETMANSKY, THOMAS ZEITZOFF. 2014. Terrorism and Voting: The Effect of Rocket Threat on Voting in Israeli Elections. *American Political Science Review* 1-17. [CrossRef]

10. Ildefonso Mendez, Isabel M. Cutillas. 2014. Has immigration affected Spanish presidential elections results?. *Journal of Population Economics* **27**, 135-171. [CrossRef]

11. Aziz Z. Huq. 2013. The Political Psychology of Counterterrorism. *Annual Review of Law and Social Science* **9**, 71-94. [CrossRef]

12. Andrew Healy, Gabriel S. Lenz. 2013. Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy. *American Journal of Political Science* n/a-n/a. [CrossRef]

13. CÉSAR GARCÍA-DÍAZ, GILMAR ZAMBRANA-CRUZ, ARJEN VAN WITTELOOSTUIJN. 2013. POLITICAL SPACES, DIMENSIONALITY DECLINE AND PARTY COMPETITION. *Advances in Complex Systems* **16**, 1350019. [CrossRef]

14. Henning Finseraas, Ola Listhaug. 2013. It can happen here: the impact of the Mumbai terror attacks on public opinion in Western Europe. *Public Choice* **156**, 213-228. [CrossRef]

15. Jose G. Montalvo. 2011. Re-examining the evidence on the electoral impact of terrorist attacks: The Spanish election of 2004. *Electoral Studies* . [CrossRef]

16. Pau Castells, Francesc Trillas. 2011. The effects of surprise political events on quoted firms: the March 2004 election in Spain. *SERIEs* . [CrossRef]