Polit Behav (2013) 35:429–452
DOI 10.1007/s11109-012-9205-1

ORIGINAL PAPER

Tactical Voting and Party Preferences: A Test of Cognitive Dissonance Theory

Jørgen Bølstad · Elias Dinas · Pedro Riera

Published online: 30 June 2012

Abstract Studying the development of stable political attitudes, political scientists have argued that repeated voting for a political party reinforces initial party preferences, in a seemingly mechanistic process of habit-formation. However, the empirical evidence is scarce and the theoretical framework underdeveloped. Does the act of voting for a party improve an individual's evaluation of this party? If so, is this effect simply due to habit-formation, or a more complex psychological mechanism? Drawing on cognitive dissonance theory, we examine the act of voting as a choice inducing dissonance reduction. We go beyond existing research, by focusing on tactical voters—a group for which the notion of habitual reinforcement does not predict an effect. The analyses reveal a positive effect of the act of voting tactically on the preferences for the parties voted for and may thus call for a revision of the traditional understanding of the role of voting in shaping party preferences.

Keywords Party preferences · Partisanship · Party identification · Cognitive dissonance · Tactical voting · Genetic matching · Multiple control groups

J. Bølstad
ETH Zurich, Zurich, Switzerland
e-mail: joergen.boelstad@eup.gess.ethz.ch

E. Dinas (✉)
Nuffield College, University of Oxford, Oxford, UK
e-mail: elias.dinas@nuffield.ox.ac.uk

P. Riera
European University Institute, Florence, Italy
e-mail: pedro.riera@eui.eu

Introduction

The central question in the study of voting behavior is how people reach their voting decisions. During the last half century, a variety of analytical models have been employed to examine this question empirically. Despite their differences, most of these models share a common feature: the assumption that vote choices reflect party preferences, which can be predicted by attitudinal indicators. This line of reasoning is clearly useful for the explanation of vote choice, but it ignores the possibility that vote choice exerts an influence on party preferences, because vote choice is seen as the ultimate stage in the process of party preference formation.

Although this non-recursive logic has prevailed in the study of electoral behavior, the possibility that electoral participation itself influences people's attitudes cannot be ruled out. Even a quick glance at the early literature of partisan development would illustrate this possibility. According to Converse (1969), party preferences reflect accumulated electoral experience: as individuals gain experience, their partisan outlooks consolidate. At least since McPhee and Ferguson's "political immunization" thesis (1962), later transferred by Butler and Stokes to the British case (1974[1969]), elections have been thought to serve as mediators of political maturation: the observed difference in attitudinal stability between young and old has been attributed to their differing level of electoral experience (Converse 1976).

The idea that voting decisions may change people's partisan orientations was tested in some early studies of party identification (Dobson and St. Angelo 1975; Knoke 1976; Markus and Converse 1979; Howell 1980, 1981). However, these studies were mostly concerned with the short-term stability of party identification, trying to compare the Michigan view of partisanship (Campbell et al. 1960) to the revisionist school of party identification (Fiorina 1981). This meant that voting decisions, along with candidate evaluations (Brody 1991), issue opinions (Jackson 1975; Franklin and Jackson 1983; Franklin 1984) and other attitudinal factors were just one of the factors that were rather inductively included in models of partisan change.[1]

Therefore, even when the results suggested that vote choice affected party evaluations, the mechanism driving these effects remained unspecified. To the extent that the mechanism through which elections may induce attitudinal change was touched upon in the literature, the main driving force of this relationship was assumed to be the acquisition of political information (McPhee and Ferguson 1962; Converse 1969; Butler and Stoker 1974). Elections were deemed to function as a funnel of political information helping the voter to get acquainted with the political parties and their leaders. Once one gets into the polling booth, casting a ballot is not

[1] The findings from these early studies provided some evidence, albeit contested (Jackson 1975), about the role of vote choice on change in party identification. Later work on the dynamics of party identification, however, challenged the idea that partisan identities respond to such short-term factors, suggesting that this seemingly volatile picture stemmed from measurement error that attenuated short-term shifts in partisan orientations (Green and Palmquist 1990, 1994). Experimental evidence on the endurance of party identification against new information about parties' ideological positions provided further evidence to the Michigan view of partisanship as a stable predisposition toward the political parties (Cowden and McDermott 2000).

important; what matters is the experience with the electoral procedure and the exposure to the campaign messages.

The main argument developed and examined in this study is that casting a ballot actually matters. In particular, the act of voting has consequences for people's attitudes toward the party opted for, making their assessment of this party more favorable. The reason for anticipating such an effect stems from the observation, well-known among cognitive psychologists, that behavior can have attitudinal implications. This seemingly paradoxical causal arrow, from behavior to attitudes, rather than the other way around, has constituted a key focal point and produced several important theoretical contributions in the field of cognitive psychology. The leading theory in this regard is that of cognitive dissonance (Festinger 1957). The idea behind this theory is that when attitudes are not consistent with actual behavior, a certain degree of discomfort will arise, driving people to change their attitudes to reduce the inconsistency. Our aim is to test whether this logic applies to voting behavior, and, more specifically, the case of voting and party preferences. We focus on a particular group of voters for whom we could actually expect a certain degree of discomfort to arise, namely, tactical voters, in a context where they are common, namely, British general elections. As tactical voters do not vote for their most preferred party, voting may involve a certain degree of dissonance. Alternative theories of why the act of voting may influence party preferences do not seem to predict that tactical voting would have the same effect as voting normally does, which is our main reason for focusing on tactical voters, asking if tactical voting improves the assessments of the parties voted for.

After providing the theoretical framework within which this question will be examined, we show in detail how focusing on tactical voters enables the identification of models in which the act of voting is used to predict changes in people's party preferences. We provide three tests, each serving a different purpose. First, we assess the effect of tactical voting on the preferences for parties voted for by comparing these to the preferences for which similar (but non-tactical) voters would have voted for if they had voted tactically. Even when we compare similar preferences from similar voters, we find that voting tactically significantly improves subsequent evaluations of the party voted for. Second, in order to better guard against potential unobserved confounders and as a way to test the robustness of these findings, we instrument tactical voting, exploiting constituency level differences in the incentives for such voting. Again, we find a positive effect. Third, we focus on within-individual differences, comparing changes in the preferences for parties tactical voters vote for to the changes in their preferences for other parties. This final comparison also reveals a positive effect of the act of voting on subsequent preferences. We discuss the implications of these findings in the conclusion.

Cognitive Dissonance and the Act of Voting

Why should we expect the act of voting to have implications for people's party evaluations? The reason, psychologists would argue, lies in the fact that any sort of

behavioral decision may imply a binding declaration of preferences. The most rigorous theoretical framework pointing to potential attitudinal ramifications of behavioral choices is Festinger's theory of cognitive dissonance (Festinger 1957). Festinger's initial formulation was based on the idea that people may unconsciously modify their attitudes to have them conform to external facts they cannot change (Festinger et al. 1956).[2] A further important contribution Festinger made was to transfer this idea to the question of how people react to their own behavior when the latter does not conform to their preferences (Festinger and Carlsmith 1959). Various studies followed Festinger's and in most cases they confirmed the presence of cognitive dissonance. For example, researchers in experimental psychology have found consistent evidence regarding the role of behavior in the reinforcement of prior attitudes, even when the first largely conforms to the second (Aronson et al. 1991). In an experiment run by Brehm (1956, see also Egan et al. 2007), subjects were given a series of presents that they were asked to rate. They were then free to choose one of them. Having done so, they were called to rate the presents again. Interestingly, people rated the present they chose higher in the second round than in the first. Conversely, they rated other items less positively than in the first round. This logic can be directly applied to the case of elections. Even when people vote for a party they clearly prefer to others, actually doing so may still lead them to think more highly of this party than before the election. In other words, even if there is no obvious conflict between behavior and attitudes, there may be a slight dissonance that can be reduced by changing attitudes so that they even more clearly favor the behavior exhibited. This is important in our case, as we focus on tactical voters, not because we do not expect an effect for other voters, but because alternative theories do not predict an effect for tactical voters.[3]

Although cognitive dissonance could be induced by various behavioral or informational stimuli, as demonstrated by the above-mentioned studies, what is of interest here is the extent to which this need for dissonance reduction is also produced by the act of voting. Evidence for this has been rather scarce, although the theory is increasingly drawing the attention of students of political behavior. Besley and Joslyn (2001) were the first to refer to the need for dissonance reduction as the mechanism generating attitudinal change after presidential elections. These authors distinguished between two types of dissonance, one stemming from the act of voting (choice-based) and the other stemming from the outcome of the election (outcome-based). By pooling various pre- to post-election panels of the American National Election Studies (ANES), they found that voters appear to register greater change in their evaluations of the candidates than non-voters. This effect was smaller among

[2] See Stone (2000) for a review of both critical and supportive essays related to the book.

[3] Cognitive dissonance has frequently been used to explain rationalization bias in studies of voting behavior, for example with regard to economic voting (Anderson et al. 2004; Evans and Andersen 2006; Anderson 2007) and partisans' placements of parties on issues where they disagree (Eiser 1992; Sherif and Hovland 1961), but see van der Brug (2001) for a null-finding on the latter topic. Plumb (1986) uses cognitive dissonance to explain bias in voters' statements regarding the timing of their changing parties before the election, while (Groenendyk, forthcoming) argues partisans counter new negative information about their party by depicting the other party in an even more negative light, thus justifying their alignment.

those who opted for the loser and who did not have a strong preference for either of the two candidates. Although their results give credence to the idea that electoral choices may reinforce prior beliefs, the selection problem was not addressed. In other words, the analysis leaves room for confounding explanations, as voters might differ from non-voters in various respects that also influence their attitude formation.

Another attempt to estimate the causal effect of the act of voting on attitudinal change has been made by Huber et al. (2009). Exploiting the near-random rotation of candidate names in New York Primary ballots, they show that being first in the ballot increases a candidate's primary vote share between 4 and 17 % points. However, they find no carry-over rank-order effect in the subsequent general election. People who voted the first-ranked candidate do not seem more likely to opt for the same candidate (irrespective of whether he is ranked first in the ballot or not) in the next election. However, given the identification strategy, this null finding is hardly surprising. Cognitive dissonance is already low when the level of discomfort is weak. For those people who simply choose the first name on the ballot, the level of interest in the election is likely to be low and the cognitive cues generated by the act of voting very small. It is even doubtful that these voters recall the name of the candidate so that they can manifest signs of cognitive dissonance. To put it differently, the local average treatment effect (LATE) in this case applies to the subgroup of the population for which cognitive dissonance is least likely to be manifested empirically.

A more elaborate testing of this hypothesis comes from Meredith (2009). In his cohort analysis, Meredith shows that young citizens of California barely eligible to vote in the 2000 Presidential Election were significantly more likely to express a political affiliation by 2004 than those barely ineligible to vote in 2000 because they had turned 18 only some days after the election. The problem with this study, however, is that, given the limited amount of information available for each date-of-birth cohort, alternative explanations related to an increase in political interest or political knowledge for those eligible to vote cannot be adequately examined.

Such competing explanations are addressed by Mullainathan and Washington (2009). The authors pool respondents from midterm ANES surveys aged 18 to 21, with those who were 18 and 19 at the time of the respective survey having been 16 and 17 in the year of the Presidential election, and thus ineligible to vote, and those aged 20 and 21 having been 18 or 19 in the year of the election and thus having had the right to vote. They compare feelings toward the incumbent president and find substantially greater polarization in respondents' preferences among those in the second group. Although Mullainathan and Washington systematically examine and eventually rule out the possibility that the observed effects are either due to age differences between the control and the treatment group or a gap in political interest generated by electoral campaigns, their results mainly refer to the effect of the act of voting on leadership evaluations. Thus, it is still not known whether the same logic applies also with regard to party preferences. This is one of the two gaps that we aim to fill. Moreover, in contrast to Mullainathan and Washington's static design, we examine the role of voting through a set of two-wave panel studies, which allows a more explicit investigation of the impact of vote choice on party preference change.

More importantly, Mullainathan and Washington cannot distinguish between different mechanisms through which the act of voting may reinforce people's prior party preferences. At least two such mechanisms have been pointed out in this literature. In the first case, in his influential work on the reconceptualization of party identification, Fiorina (1981) found a nontrivial level of continuity in voting patterns among partisans that made him conclude that there is an inertial feature of party identification, which makes vote choice a standing decision, or habit—but how does vote choice become a standing decision? Fiorina argues that this is the result of repeated vote choice for the party one really prefers and initially opts for on rational grounds. From this perspective, the memory of a previous voting decision has important implications for one's current party evaluations. This relationship, however, is not due to a dissonance reduction mechanism. Rather, vote choice works as a self-reinforcing shortcut of retrospective evaluations, what he calls a "running tally."

The second formulation of a mechanism is found in the work of Grofman et al. (2009). Exploring the relationship between partisanship and vote choice, Grofman et al. (2009, p. 67) employed a political scientists' analogy of the saying: "you become what you eat:" individuals become aware of their attitudes by observing their own behavior (Grofman et al. 2009, p. 68). Party preferences, according to this logic, are gradually strengthened through a process that resembles Bem's self-perception theory (Bem 1972). Self-perception is more parsimonious than cognitive dissonance in that it postulates the same reinforcing effect of behavior on attitudes without referring to any sort of discomfort arising from a conflict between the two. According to the originator of the theory, individuals come to know their own attitudes, emotions and other internal states partially by inferring them from observations of their own overt behavior or the circumstances in which this behavior occurs. "[When] internal cues are weak, ambiguous, or uninterruptable, the individual is […] in the same position as the outside observer" (Bem 1972, p. 2).[4]

It seems that in order to differentiate between cognitive dissonance and these alternative mechanisms, one needs to test the effect of the act of voting in a setting in which actual choice and true party preference do not coincide. In such a context, any effect of the act of voting could more convincingly be attributed to the mechanism of dissonance reduction anticipated by Festinger rather than to a more mechanistic process of partisan anchoring advocated either by the "running tally" or the "self-perception" hypotheses. In all previous studies, people are assumed to vote for their most preferred party. Although selection problems are sometimes addressed in these studies, it is only by weakening the link between party preference and vote choice we can assess the role of dissonance reduction as mechanism enhancing prior party evaluations.

[4] Another theory of attitudinal change, balance theory, could also be of interest here, because it anticipates that people try to maintain balance between their attitudes and preferences. We do not refer to this theory in the main text because it does not directly relate to the impact of behavior. Balance theory is a more general theory of how people update their perceptions so as to preserve balance in light of new information about the stimulus of interest (Heider 1958).

Design

To identify the effect of vote choice on party preference change in a context in which we can find conflict between attitudes and behavior, we focus on voters who do not opt for their first preference. By definition, one such group is tactical voters, i.e. people who opt for a second (or lower) preference in order to avoid having a less preferred party win the election (Fisher 2004).[5] If cognitive dissonance is one of the mechanisms behind the effect of the act of voting, tactical voting should also create attitudinal change. In contrast, if people become locked into voting patterns through a self-perception or a habit formation process, such strategic considerations, not based on the repetition of the same choice for the party actually preferred, should not contribute to this pattern. Thus, the hypothesis tested here is that people who opt for a party on tactical grounds will end up liking the party more after having voted for it. We need to emphasize, however, that a positive finding in this case does not imply that voting sincerely does not also increase the level of party sympathy—in fact, that is to be expected based on earlier studies. We believe finding evidence in favor of cognitive dissonance in the present case makes it more plausible that this mechanism is also present for other voters. However, such evidence does not rule out the possibility that vote choice also affects post-election attitudes through a different mechanism, such as a habit formation or self-perception process. The problem is that several mechanisms may explain the effect in the case of sincere voting, which is why we prefer to focus on tactical voting, where the case for an interpretation based on cognitive dissonance is clearer. Moreover, as will be also illustrated in the next section, it is much harder to address problems of selection into the treatment condition—voting for a party—among sincere voters rather than among tactical voters. Shifting our attention to the latter, we are more confident that our findings are not driven by unobserved heterogeneity.

We focus on the UK because of the considerable amount of tactical voting produced by its electoral system.[6] For our dependent variable, we seek panel data that provide voters' assessments of the political parties before and after elections.

[5] In our view, there are at least two other situations in which some sort of discomfort can be expected to emerge in the context of democratic elections: first, in a two-round majority runoff system, when the party/candidate preferred by the voter is eliminated in the first round; and second, in all types of electoral systems, when the party/candidate preferred by the voter strategically abstains from entering the race because they have zero chances of obtaining a seat/winning the election. Compared to these two groups, tactical voters are preferable because it is easier to rationalize—and thus creates less of a need to reduce dissonance—a vote for a party different than one's most preferred party when the latter does not participate in the election than when it takes part but bears little chance of winning in the constituency.

[6] The single-member district and plurality rule used in Britain is often seen as the archetypal context within which tactical voting is likely to flourish. Given that only one candidate is elected from each district under this system, voters often have incentives to switch from their first preferences to candidates who have better chances of beating a less favored candidate. Although the exact way to measure this kind of voting behavior is still relatively disputed in the existing literature (compare for example Niemi et al. 1993 with Fisher 2004), aggregate-level studies on the basis of constituency vote returns (Shively 1970; Spafford 1972; Johnston 1981; Curtice and Steed 1988; Johnston and Pattie 1991; Evans et al. 1998; Kiewiet 2009) and studies using survey data (Cain 1978; Heath et al. 1991; Evans 1994; Alvarez and Nagler 2000; Fisher 2004; Fisher and Curtice 2006) alike have established that tactical voting is a stable feature of British elections.

For the UK, such data are available from the British Election Study for the elections of 2001 and 2005. In order to maximize the number of potential observations, we have pooled the data from these two elections, for which the survey items needed for our analysis have an identical wording format and the same response categories. The most important indicator for our design comes from the following question: "On a scale that runs from 0 to 10, where 0 means strongly dislike and 10 means strongly like, how do you feel about [the Labour Party]?" We are interested in the changes in this variable from pre- to post-election.[7] In 2001, this question was asked across the whole UK for the Labour Party, the Conservatives, and the Liberal Democrats. In addition, it was asked for the Scottish National Party in Scotland, and for Plaid Cymru in Wales. In 2005, the Green Party was also included for the whole UK, and the United Kingdom Independence Party for England.

It needs to be stressed that our focus is not simply on voters, but on their preferences. For a single voter, several preferences are registered in our data (one for each party mentioned above). We are more interested in some preferences than others, especially those tactical voters vote for. We therefore transform our dataset from one in which the observations are at the individual-level, to one in which observations are at the more finely grained party*individual-level. For each individual, we have as many observations as there were parties covered by the respective survey.

The selection of specific preferences for our analyses creates a need for shorthand names for the preferences we select. Accurate labels would otherwise require about ten words for our main groups of interest, and thus not be very practical. Although this is not an experimental study, our aim is to emulate an experiment as closely as possible, minimizing the problems resulting from not being able to assign vote choice randomly. For these reasons, we find it useful to adopt a bit of experimental terminology and refer to the observations for which we expect an effect as treated, and those to which we compare them as part of a control group. We want these groups to be similar in all important respects, just as if the treatment had indeed been randomly assigned. To the extent we achieve this, we can compare the groups, and assess whether there is an effect as hypothesized. As a way to crossvalidate our results, we engage in multiple control group testing (Daniel et al. 2008; Lu and Rosenbaum 2004; Rosenbaum 2010).

As mentioned, we are especially interested in tactical voters. We identify voters who make a conscious choice to vote tactically, using a question regarding the reasons for vote choice: "People give different reasons for why they vote for one party rather than another. Which of the following best describes your reasons?" If respondents choose "I really preferred another party but it stood no chance of winning in my constituency," or if they voluntarily declare "I voted tactically," we count them as tactical voters.[8] However, we exclude "false" tactical voters from this group. That is, those who on the follow-up question regarding which party they

[7] In 2001, the election was held on June 7, the pre-election wave was conducted from March 3 to May 14, and the post-election wave started immediately after the election, on June 8, and lasted until July 30. In 2005, the election was held on May 5, the pre-election wave was conducted from February to April 12, and again the post election wave started immediately after the election, on May 6, and lasted until July 4.

[8] The other alternatives are: "The party had the best policies," and "The party had the best leader."

really preferred reported the same party as the one they voted for. We have also excluded as false tactical voters those who voted for a party they preferred above all others, as measured by the pre-election party preferences mentioned above. Thus, we require that tactical voters make a conscious choice to vote tactically, and that they like another party at least as well as the one voted for.[9] The *treatment group*, then, consists of the preferences for the parties these tactical voters vote for.[10] This means that in our treatment group there is only one case per individual, i.e. only one party*individual combination is chosen, the one that corresponds to the party voted on tactical grounds. We will conduct three analyses using this treatment group, but based on two different control groups, and we therefore introduce these control groups in relation to the specific analyses.

To minimize potential bias in estimating our effects, we select a set of covariates that appear likely to capture the most plausible confounders. We will use these both as controls in our analyses below, and in our assessment of the treatment and control groups. One important factor is the level of interest in the election, as tactical voters may be more (or less) interested than others, and they may like parties better after having paid attention to their campaigns. We will include a measure of such interest, based on a question in the pre-election survey: "How interested are you in the general election that is likely to be held soon: 1 very interested, 2 somewhat interested, 3 not very interested, or 4 not at all interested?''[11] Moreover, we introduce a self-reported scale (0–10) of attention paid to politics as well as two indicators of campaign exposure, based on the following questions (including follow-ups that identify specific parties): "Did a canvasser from any party call at your home to talk with you during the election campaign?;" "Did anyone from a political party telephone you during the election campaign to ask you how you might vote?;" and "Did you see any of the Party Election Broadcasts that were shown on television during the election campaign?" We combine the indicators for having been canvassed or phoned by a political party into one, because so few respondents were phoned by political parties during the campaign that the indicator is ill-suited for matching.[12] These are, then, binary indicators of campaign exposure, at the party*individual-level, varying within individuals. We also include age, gender, and the age at which education was ended.[13]

[9] See Fisher (2004) for a theoretical and empirical justification of this measurement strategy.

[10] This may, however, involve a risk of excluding the voters experiencing the most cognitive dissonance. If a voter initially prefers a small third party, but ends up voting for a one of two larger parties, cognitive dissonance may be so strong as to make this voter believe she even preferred the chosen party before the election. If she thus fails to state that hers was a tactical vote, we will not identify her as a tactical voter, and not register this effect. Another potential problem is that individuals may over-report having voted for the winner, but it seems less likely this would influence our results, and, in both cases, the potential bias would most likely serve to conceal rather than exaggerate the effect we seek to identify.

[11] We have reversed the number assigned to these alternatives, however, in order to ease interpretation.

[12] The alternative would be dropping this variable from the analysis, which we find less attractive. The new indicator is 1 if either of the two initial indicators are 1, and 0 otherwise.

[13] Age is included as a continuous measure; for the gender dummy, 1 is male; education is coded as follows: four dummies have been constructed corresponding to the following categories: finished education at the age of 16 or lower; finished education at the age of 17 or 18; finished education at the age of 19 or higher; still in education. The first of these categories serves as the reference category in the parametric analysis.

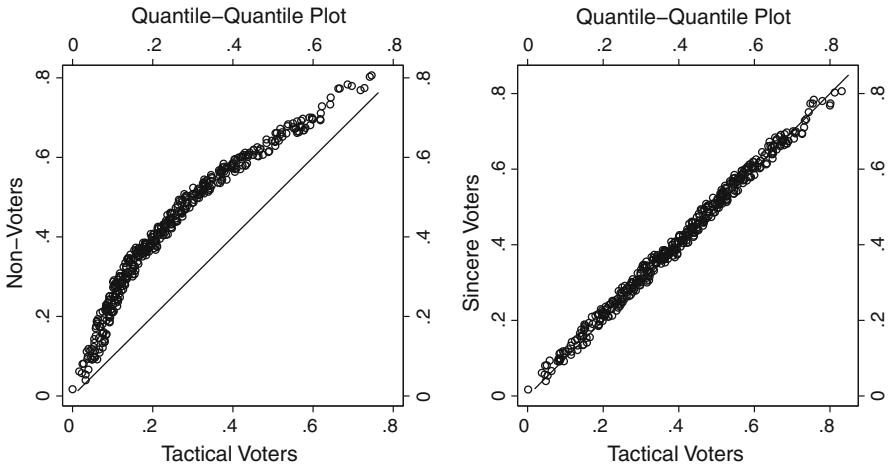

Fig. 1 Quantile-quantile plots of the predicted probability of voting tactically for tactical voters versus non-voters (*left*) and sincere voters (*right*). Note: The predictions are based on a logistic regression of tactical voting (vs. non-voting) on all covariates listed in the text

Analysis I

The remaining question is that of selecting a control group. As our discussion of the literature illustrates, the few studies that have aimed to test the presence of cognitive dissonance in the effect of vote choice, have mostly compared sincere voters to non-voters. However, important differences in terms of political interest, prior party evaluations and other potentially important predictors of both vote choice and posterior party evaluations imply that such a comparison is likely to produce biased estimates. We do not believe non-voters constitute the best available group with which to compare either sincere or tactical voters. We want a control group that is as similar to the treated group as possible, in order to rule out alternative explanations relating to their differences. Non-voters do not only differ from tactical voters on the tactical aspect of voting, but also on the turnout aspect. Thus, we want to select our control group among the preferences of respondents who actually voted in the elections.

Figure 1 illustrates the differences between the groups in question. Our expectation that tactical voters are more similar to other (sincere) voters than to non-voters holds empirically. The first graph presents a quantile-quantile plot of the predicted probability of voting tactically, based on a logistic regression on all covariates mentioned above, among tactical voters versus non-voters.[14] The plot on the right is similar, but it focuses on tactical versus sincere voters, instead of non-voters.[15] The two plots clearly illustrate that tactical voters are more similar to

[14] The non-voters have been identified by the question: "Talking with people about the general election […], we have found that a lot of people didn't manage to vote. How about you, did you manage to vote in the general election?''

[15] The predictions for both plots are based on the regression of tactical voting versus non-voting, because regular voters are so similar to tactical voters, a regression based on them hardly yields any significant

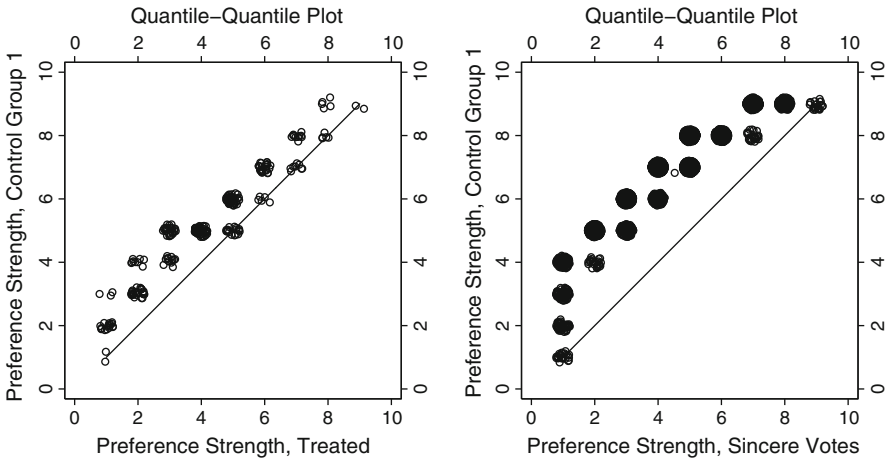

Fig. 2 Quantile-quantile plots of the pre-election preferences voted for by tactical voters (treated) and those voted for by sincere voters against the other preferences of sincere voters (control group 1)

sincere voters than to non-voters. In fact, tactical and sincere voters appear virtually identical. Thus, by selecting our control group among the preferences of sincere voters instead of those of non-voters, we will be comparing preferences from individuals that are very similar in important respects.

It is also worth looking at imbalances in terms of the initial levels of our outcome variable, party preference strength. Figure 2 presents two quantile-quantile plots. The first shows the distribution of our treated preferences (those voted for by tactical voters), against those not voted for by sincere voters. The second plot shows the same, but for the preferences voted for by sincere voters instead of those of tactical voters. Compared to the pre-election levels of the preferences voted for by sincere voters, the equivalent preferences for tactical voters are much closer to those in our potential control group (those not voted for by sincere voters). We are thus confident that, apart from the theoretical motivation for focusing on tactical voters, we also gain on methodological grounds, since we focus on a treatment group that is more similar to our control group than that which has most frequently used in previous studies.

As shown in Fig. 1, tactical and sincere voters are almost perfectly balanced in terms of their most relevant characteristics. However, we want to go further and also select the most relevant among the preferences of the sincere voters. More specifically, we want the preferences they would have voted for if they had voted tactically. As mentioned, as a first step, we select all the preferences sincere voters could have possibly voted for if they had voted tactically (i.e. not the preference for the party actually voted for), but we want to be more specific than this. A look at the data

Footnote 15 continued
estimates. This means that the plotted probabilities of tactical voting in the plot on the right are higher than they would have been if they were predicted based on a regression including sincere voters (which also makes the two plots more comparable), but the important point to notice is the balance between tactical and sincere voters, which would be evident regardless of which regression was used to produce the predictions.

suggests more than 75 % of tactical voters vote for their second most preferred party, but this also means that about 25 % do not. In order to select the right preferences as controls, we construct a variable measuring the rank of an individual's preferences. In other words, their strongest preference has rank 1, their second strongest rank 2, and so on. When two preferences are tied (e.g. they are both rank 1), the next preference still gets the rank it would have gotten if there were no tie (in this example, 3).[16] As explained below, this measure of preference rank is used in our matching scheme, to obtain optimal balance between treated and controls.[17] Thus, as comparisons for the preferences voted for by tactical voters, we will focus on the equally ranked (but not voted for) preferences of sincere voters, and we will refer to them as *control group 1*.[18]

We use genetic matching as a method to match respondents in terms of relevant covariates (Diamond and Sekhon 2008).[19] We prefer matching because it is non-parametric (thus, not dependent on functional form assumptions) and makes the analysis more transparent (Eggers and Hainmueller 2009, p. 9). In our matching procedure, we include pre-election data on the dependent variable in a fully factorized fashion, including a dummy for each possible value, in order to avoid functional form assumptions with regard to how this variable affects the outcome (see Arulampalam et al. 2000 for a discussion of the intial conditions problem). This

[16] Further, to make the variable work in a logical manner, when tactical voters vote for a party that is tied with another (say, both were rank 1), the tactically voted preference is ranked below the other (in this case, 2), as the fact that the voter sees the vote as tactical implies the other party was actually preferred. According to a similar logic, for sincere voters, when the party voted for is tied, the other party is ranked below the party voted for, as the vote implies that this other party was less attractive.

[17] To further ensure that we focus on the right preferences, we also exclude from the analysis preferences that are zero (minimum) before the election. Such preferences indicate that parties are "strongly disliked,'' which means they are extremely unlikely to form the basis of a vote, whether tactical or non-tactical. In other words, few or none of the treated preferences are expected to be of this kind, while some are to be expected among the potential control preferences. This would be a source of imbalance, and as such preferences would be of little interest in the analysis, they are excluded from both the treated and the control group. We also exclude preferences that are ten (maximum) before the election. The reason it is possible for treated preferences to be ten is that we have allowed tactical voters to vote for a party tied with their first preference. However, such tactical voters are not the most interesting ones, especially when they vote for a party for which they express a maximum preference before the election—suggesting that they might as well consider their vote non-tactical. Thus, again having as our motivation to maximize balance—there is no point in which the dependent variable has been taken into any consideration thus far—we focus our analysis on relevant preferences, excluding from both the treated and the control group those that have a value of ten before the election.

[18] We could use only one party preference for each individual. However, to supply more cases that could be used as controls in the matching procedure, in our estimation presented in the next section, we have allowed each individual in the group of sincere voters to be represented with more than one party*individual combination. For instance, a sincere voter may supply two observations in the control group, one related to her second and one related to her third preference. Again, the motivation for this decision is to facilitate the process of finding balance between the control and the treated groups in terms our set of observed pre-treatment characteristics. We have also tested the sensitivity of our results, by only selecting one party*individual combination for each individual in our control group is chosen. The effects are almost identical to those presented here.

[19] Genetic matching is a generalization of propensity score and Mahalanobis distance matching that uses an evolutionary search algorithm to determine the weight given to each baseline covariate (Sekhon and Mebane 1998). Genetic matching has shown to outperform in terms of Mean Squared Error alternative matching methods both when the Equal Percent Bias Reduction assumption holds and when it does not. In the latter case, it also performs much better in terms of bias (Diamond and Sekhon 2008).

set of dummies is referred to in the figures below as "Initial Preference Strength: [1–9]." As mentioned, we also match on the rank of each preference, which we also include fully factorized and refer to as "Initial Preference Rank: [2–4]." Apart from these variables, the full set of covariates used in the matching consists of the two campaign-related variables, interest in the election, attention to politics, age, education and gender. We believe these capture the most important aspects in which our treated group might differ from control group 1.

Figure 3 shows that, while the two groups initially differ in several respects, after matching, these differences are almost perfectly removed. Before matching, the mean and distributional differences between the two groups were significant for a majority of the covariates. After matching, the lowest (t) p-value is 0.57, while the lowest Kolmogorov–Smirnov p-value is .44. For all but one of the other covariates, the p-value is 1. To evaluate the overall balance between controls and treated units, we employ an omnibus test, which captures all covariates and indicates whether the matched data as a whole appear unbalanced. We use the test proposed by Hansen and Bowers (2008), which has been shown to perform better than other summarizing tests of balance, such as those based on standardized differences or goodness-of-fit comparisons from logistic regressions (Hansen and Bowers 2008, pp. 3–6).[20] The result from this test confirms the individual balance diagnostics shown in Fig. 3. The χ^2 statistic is 0.417, which, with 18 degrees of freedom, gives a p value 1. This means it is highly unlikely that any differences between the two groups remaining after matching are different from those to be expected from a random assignment of tactical voting. Having achieved balance through our matching procedure, we estimate the LATE of voting tactically on the preferences that are or would have been voted for on tactical grounds.[21] Table 1 presents the results and shows the LATE is 0.853, with a standard error of 0.138. Thus, controlling for these covariates, we find a clearly significant effect, by which the act of voting tactically improves the assessments of the party voted for by .85 on a 0–10 scale.

Analysis II

A potential concern with the comparison between tactical voters and sincere voters is that the first group may differ from the second in unobserved characteristics that

[20] This test has some desirable properties: it culminates in a single test statistic and p value; it is based on an χ^2 approximation which seems to work well with small samples; and it appraises balance not only on the set of covariates listed in Fig. 3, but also on all linear combinations of them (Hansen and Bowers 2008). To implement the test we used the Randomization Inference tools package (RItools) in R (Hansen and Bowers 2008).

[21] We estimate the LATE here rather than the ATE because we essentially have focused our analysis on a subset of units whose potential outcomes are affected by the treatment assignment mechanism, i.e. voted preferences of tactical voters and similar, but non-voted preferences of sincere voters. In other words, we estimate a LATE, defined by the selection of our treatment and control groups. This is analogous to a regression discontinuity (RD) design, although we do not use such a design here. In RD, only observations around a given threshold are examined because cases far from this cut-off point are deemed to differ from those near the threshold in various observable and unobservable ways. Here, this applies for sincere voters' party preferences, among which we only select those most similar to the treated preferences.

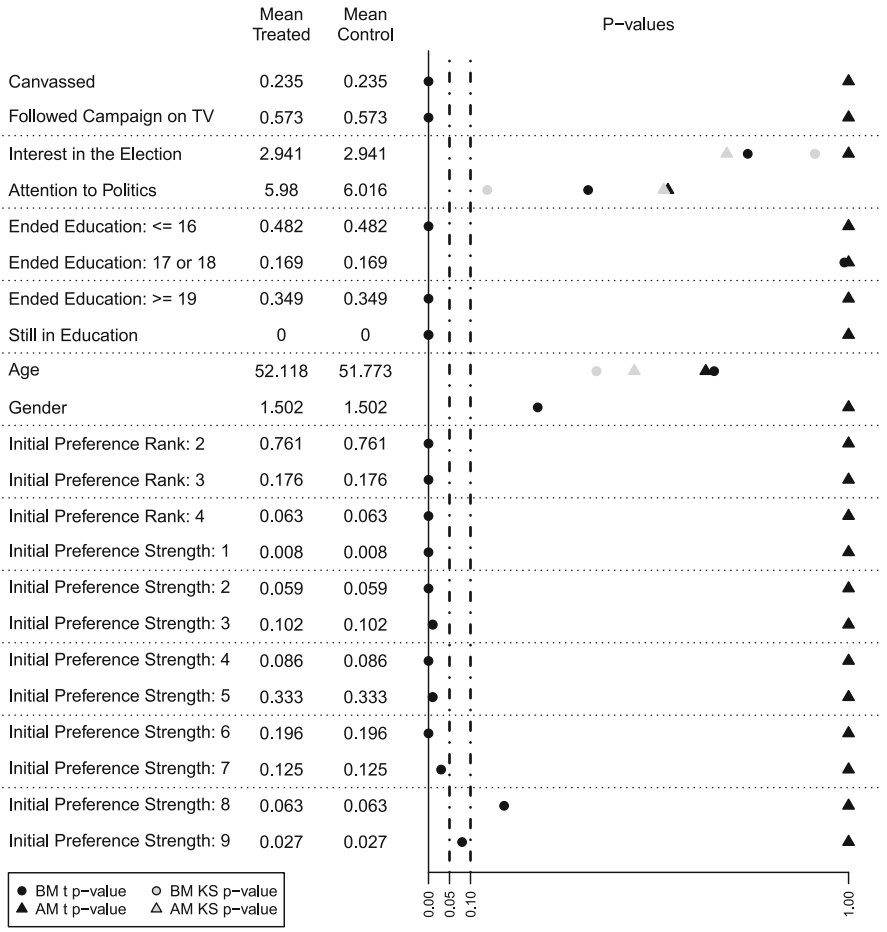

Fig. 3 Improved balance from genetic matching in the case of treatment versus control group 1

Table 1 LATE of tactical voting based on control group 1

Effect of voting tactically on preferences that are or would have been voted for on tactical grounds.	
Local average treatment effect (LATE)	0.853 (0.138)

Abadie and Imbens' standard errors are reported in *parentheses* (Abadie and Imbens 2006)

are not captured by the set of controls included in the analysis. To be more confident that our results are not a product of such differences, we conduct a further analysis. We build on the previous analysis, in which we matched the preferences voted for by tactical voters to those sincere voters would have voted for, if they had voted tactically. Within the matched dataset, we now instrument tactical voting using a constituency-level context variable, aiming to isolate variation in tactical voting that is not due to unobserved individual differences. Our analysis is motivated by the following question: if we select voters who are similar in all important observed

respects, but some of which live in constituencies that cause tactical voting and others do not, how does this influence their preferences?

We exploit a contextual feature that is likely to induce tactical behavior, namely the structure of parties' vote shares at the constituency level. Even when individual-level characteristics are conductive to placing a tactical vote, such a vote might not manifest itself if the constituency-level structure of vote shares does not give incentives for this type of voting behavior. Previous studies have shown that tactical voting is more likely to occur when there is a close race between the two largest parties (Evans and Heath 1993; Fisher 2002). Under such circumstances, people may vote for a party that is not their first choice in order to avoid the least liked party gaining the seat. If the closeness of the race, which we will refer to as marginality, thus predicts tactical voting, it may work as a district-level instrument. We use a binary variable to identify marginal constituencies, i.e. districts in which the difference between the first and the second party was less than 5 % points in the previous election (1997 for 2001-respondents and 2001 for 2005-respondents). The expectation is that we find more tactical voters in these constituencies than in those with safe seats. Using this as an instrumental variable will give us a particular LATE, namely the effect for those who respond to the instrument (known as "compliers," which in this case means individuals who would vote tactically if living in a marginal district and would not do so if living in a non-marginal district).[22]

To provide an unbiased estimate of the effect, the marginality variable needs to satisfy some criteria. These have already been discussed in the casual inference literature, and we will therefore only briefly summarize them here (see Angrist et al. 1996 and Kern and Hainmueller 2009 for a formal and an intuitive explanation of each of these criteria, respectively). The criteria are: (a) Independence: people have not chosen to live in or move from constituencies where the level of marginality is low. (b) Exclusion: people living in districts with high levels of marginality do not differ from those with lower values in other terms that might be relevant for our outcome variable. (c) First stage: information about level of marginality helps predicting tactical voting. (d) Monotonicity: there are no defiers, i.e. no people that would have been discouraged from voting tactically if the constituency characteristics were favorable and would be encouraged to do so if the constituency characteristics were not favorable for tactical voting.

Looking at these assumptions in the reverse order, the monotonicity assumption is not very problematic, as it appears quite unlikely that people would have voted tactically had they resided in a non-marginal constituency, where the chances of influencing the outcome are much lower, whereas they would not have cast a tactical ballot had they resided in a marginal constituency. Assumption (c) can be tested empirically. As is shown in the first column of Table 2, knowing whether a respondent resides in a marginal district helps us predict whether she or he will vote tactically. The estimate comes from a logistic regression and hence it represents the log odds of voting tactically as a result of living in a constituency with a contested

[22] Note, however, that the LATE we identify is further complicated by the fact that we use the matched dataset, which already focuses the analysis on a particular subset of the population.

Table 2 IV estimation of the effect of tactical voting on party preferences

	2SLS (first stage)	2SLS (second stage)	LARF (second stage)	Accounting for cluster assignment
First/second stage estimate	1.04	2.27	2.26	.094
	(.37)	(1.17)	(1.18)	[.037–.150]
N	438	438	438	141

Note: In the first column are log-odds from a logistic regression of the treatment on marginality, controlling for all covariates described in the text. Standard errors are given in *parentheses*. The *second* and *third columns* report the 2SLS and LARF estimates respectively. The *fourth column* shows the results from the randomization-based method of Hansen and Bowers (2009). More details about the estimation procedure are given in footnote 24

seat. Even when a long series of contextual and individual characteristics are included in the estimation (not shown in the table but described in detail below) to address concerns with regard to the last two assumptions discussed below, marginality appears to increase the likelihood of voting tactically.

Exclusion (b) implies that marginality does not influence party preferences after an election in any other way than through its effect on the propensity to vote tactically. This assumption would be violated if, for example, parties devote more campaign resources in marginal districts, and marginality exerts a direct effect on party preferences through this mechanism instead of tactical voting (Grimmer et al. 2011). To address this issue, we include a series of individual- and district-level covariates. First, although we stick to the already matched dataset, we condition on all variables included in the matching analysis. Some of these variables are party-specific and explicitly designed to capture campaign dynamics, as they refer to party canvassing, interest in the election, and watching the leaders' debate. In addition, we include the following party-specific variables: (1) the vote share of the three major parties, Labour, the Conservatives and the Liberal Democrats, in that election; (2) the vote share of the third largest party in the constituency, which may not always be one of these three parties; (3) the level of turnout in the current election. Finally, we control for pre-election party preferences in a fully factorized fashion. By conditioning on these variables, we believe there is limited room for violation of the exclusion restriction.

The assumption of ignorability (a) requires that our instrument is as good as randomly assigned. Again, it is difficult to sustain this claim without conditioning on a vector of covariates (Grimmer et al. 2011). We have therefore used information available at the district level on various sociodemographic characteristics. In particular, the following variables are also included as controls: the proportion of the population in the constituency who are aged between 16 and 24, between 25 and 39, and between 40 and 64, as well as the proportion of: non-whites, unskilled workers, semi-skilled workers, managers and professionals, technicians, economically non-active adults, and owner occupied households. Education, age and gender differences are already captured with our individual-level covariates. We believe the controls used to satisfy exclusion and ignorability together ensure marginality only

exerts an effect on party preferences through its influence on voters' propensity to cast tactical rather than sincere votes.

Turning to estimation issues, we confront two problems. First, the typical 2-stage-least-squares estimator (2SLS) only reduces to the Wald estimator and gives us the LATE in the absence of covariates. When covariates are included, 2SLS no longer identifies the LATE, unless we assume constant treatment effects (Kern and Hainmueller 2009, p. 389). In addition to the 2SLS estimator, we therefore use an estimation method that allows for heterogeneous treatment effects, namely Abadie's Local Average Response Functions (LARF, Abadie 2003).[23] As can be seen from the second and the third column of Table 2, which present the results from the 2SLS and the LARF estimations respectively, the effect of tactical voting on party preferences, when instrumented through the level of marginality, is quite similar regardless of the chosen method of estimation.

This is important because it justifies the decision to assume homogenous treatment effects when addressing the second challenge for our estimation, namely the fact that our instrument is assigned at the district-level whereas treatment status is assigned at the individual*party-level. Clustered assignments treated as individual assignments can jeopardize inference, as the recent discussion over the effectiveness of telephone messages on Get-Out-The-Vote experiments has shown (Imai 2005; Gerber and Green 2000, 2005; Hansen and Bowers 2009). To address this problem, we employ a method proposed by Hansen and Bowers (2009), which addresses the problem of cluster assignment without losing information as a result of aggregating the data at the cluster level. This method draws on randomization inference and performs well when the control group is relatively large, as is the case here, given that 90 percent of the constituencies are coded as non-marginal. Because this method is primarily designed for binary outcomes, we recode our dependent variable, which now becomes a dummy denoting preference scores higher than 5, the median value in the original measure. Another implication of applying this method is that we only capture the intention-to-treat effect, i.e. the effect of marginality without accounting for imperfect compliance.

To briefly describe the method, the first step is to estimate a logistic regression of this variable on all control variables, using only the non-marginal constituencies. The first step is thus using all information both at the individual and the district level. The second step is to use the model estimated in the first step to generate fitted values for both treatment and control units. From the fitted values we then obtain residuals, by subtracting the fitted values from the observed data. We then aggregate these residuals at the constituency level, which is the cluster level in our case. These residuals represent the observed proportion of party preferences higher than 5 minus the proportion we would predict if there were no marginal constituencies, i.e. if all subjects—which are now constituencies—were withheld from the treatment assignment condition. The average of these residuals constitutes the average effect attributable to the treatment (Hansen and Bowers 2009).[24]

[23] The LARF estimation has been implemented by using Abadie's cls.m function for linear response in Matlab: http://www.hks.harvard.edu/fs/aabadie/cls.m.

[24] To be more precise, let us summarize the estimation procedure by following the notation of Hansen and Bowers (2009). Let r_i denote constituency i's observed proportion of party preferences higher than 5 and let C denote the group of non-marginal constituencies, so that r_i denotes the same proportion in the

The results appear in Table 2. As shown in the second and third column, the 2SLS and the LARF estimators agree that voting tactically improves the assessments of the parties voted for by approximately 2.3 points on a 0–10 scale. Hence, the magnitude of the found effect is considerable. Moreover, the assumption of homogenous treatment effects does not seem to be severely violated in this case, as the LARF estimator generates virtually the same result as the 2SLS estimator. Finally, the last column of the table accounts for the fact that instrument and treatment are assigned at different levels. Following the randomization-based inference method proposed by Hansen and Bowers (2009), we find that the anticipation of a close race increases the proportion of high party preferences in a given constituency by almost 10 per cent, an estimate that ranges between 4 and 15 percent.

Analysis III

As indicated above, we also use a different control group to make sure the results in Analysis I are not a product of unobserved differences between our treatment and control group 1. In this analysis, we conduct a comparison in which individual characteristics are rendered irrelevant. We use the preferences not voted for by tactical voters as a control group, while the treatment group still consists of the preferences voted for by the same individuals. This provides an even better set-up to isolate the impact of the act of voting from potential confounds stemming from differences in individual characteristics, such as levels of political interest. If changes in the preferences of tactical voters are due to some particular characteristics that they do not share with sincere voters, this would not bias our second comparison. Our new control group, henceforth *control group 2*, consists of the three first preferences of tactical voters (excluding the treated ones, i.e. those for parties voted for). The plot in Fig. 4 below is a quantile-quantile plot of the pre-election strength of the treated preferences (those for parties voted for by tactical voters) against those of control group 2, and demonstrates excellent balance in terms of these (which is one of the few ways in which there could be imbalance in this case). What it shows is that, before the election, tactical voters' preferences for the party eventually opted for have a distribution very similar to that of their other three strongest preferences.

We use the same analytical approach here as in Analysis I, but this time using control group 2, comparing the preferences voted for by tactical voters to those not voted for by the same group. In this analysis we leave the individual-level covariates

Footnote 24 continued
absence of marginal constituencies. The observed outcome can be written $\sum_U r_i$, where U denotes all constituencies, while the proportion of high party preferences one would extrapolate only from the group of non-marginal seats can be written $\sum_{i \in U} \hat{r}_{ci}(\hat{\beta})$, where $\hat{\beta}$ represents the logistic regression coefficients from the logistic regression of the outcome on the covariates among the control group. The final estimate comes from the difference between $\sum_U r_i$ and $\sum_{i \in U} \hat{r}_{ci}(\hat{\beta})$. The 95 % confidence bands result from a calculation of the variance of $\bar{\hat{r}}_C$, i.e. $\hat{V}(\bar{r}_C) = (1 - n/N)s^2[r_i : i \in C]/n$, where $n = |C|$, $N = |U|$ and $s^2[(r_1, \ldots, r_J]) = (J-1)^{-1} \sum_1^J (r_j - \bar{r})^2$.

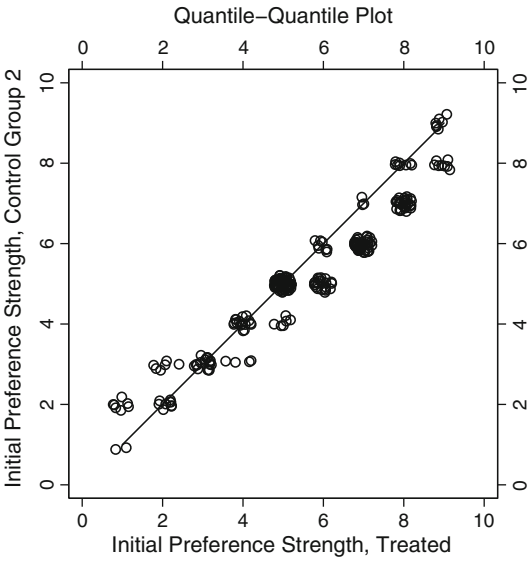

Fig. 4 Quantile-quantile plot of the pre-election preferences voted for by tactical voters (treated) against the other preferences of tactical voters (control group 2)

out, as they are of little relevance for this within-individual comparison and as limiting the number of covariates tends to improve the matching on those covariates that are included (Robinson et al. 2009, p. 354). The only covariates that could possibly confound the effect of the act of voting are the two party-specific campaign-exposure dummies described above. Thus, we match only on these two along with the pre-election strength of the party preferences. Figure 5 shows the improvement of balance resulting from this procedure. Although in this case balance is not achieved as consistently as in the previous analysis, the lowest p value (0.16) is still far from the 0.05 and 0.1 conventional limits. The χ^2 statistic of the overall test of balance between the matched treated and control groups is 0.258, which, with 10 degrees of freedom, gives a p value of 1 (cf. Hansen and Bowers 2008). Table 3 reports the results, showing that the LATE is 0.556. This estimate is fairly similar to the one reported in Table 1, which increases our confidence in the results. These results suggest that individuals who choose a party on tactical grounds on average become half a point more favorable in their assessments of this party on a 0–10 scale (as compared to their other preferences of similar strength).

Conclusion

The aim of this study was to examine the effect of the act of voting on subsequent party preferences, distinguishing it from the effect of prior preferences and other plausible confounders. While cognitive dissonance theory suggests the act of voting has an effect in itself, the traditional understanding in political science suggests that the role of voting is merely one of strengthening initial preferences that produce a certain vote choice, thus reinforcing existing inclinations. Prior studies are plagued by the problem that these theories yield similar predictions and that initial

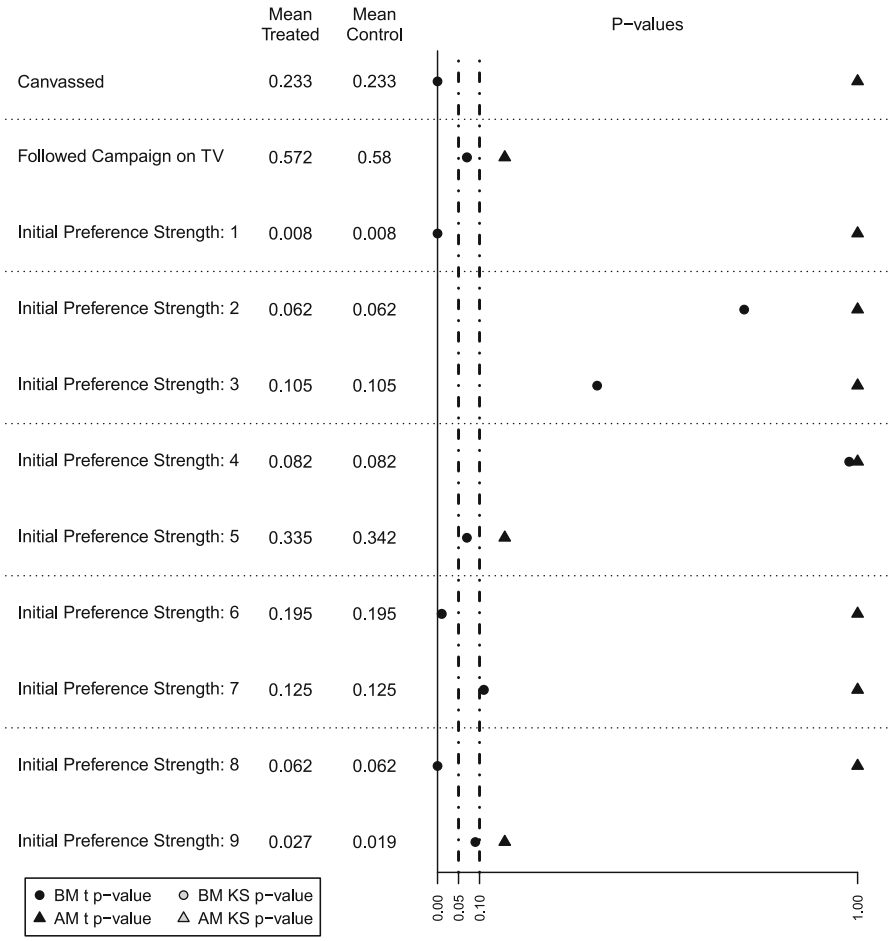

Fig. 5 Improved balance from genetic matching in the case of treatment versus control group 2

Table 3 LATE of tactical voting based on control group 2

Effect of voting tactically on the preferences of tactical voters.	
Local average treatment effect	0.556 (0.123)

Note: Abadie and Imbens' standard errors are reported in *parentheses* (Abadie and Imbens 2006)

preferences are bound to influence both vote choice and subsequent preferences. The contribution of this study therefore lies in its focus on tactical voters—a group for which only the theory of cognitive dissonance predicts an effect, as these voters do not vote for their most preferred party. Prior evidence on the presence of post-selection bias (Egan et al. 2007) suggests that Festinger's theory can be applied also to cases where the level of dissonance is very low, as when behavioral choices match initial preferences. Thus, the theory may also apply to non-tactical voters, and

a positive finding in our case makes this more plausible, but, given our focus on tactical voters, we can only conclude with regard to these.

The analysis does indeed demonstrate an effect in the case of tactical voting. This effect is significant whether our comparison is based on the relevant preferences of other voters, or those of tactical voters themselves. It is also significant when we instrument tactical voting using the constituency-level characteristic of district marginality. This means we can largely rule out individual characteristics as confounders. The effect is robust to our best efforts to control for plausible confounders. The magnitude of the effect is also notable. The average voter gives her most liked party a preference score of about 6.5, and the second most liked a score of 5. Thus, our estimates, showing that voting for a party can increase its score by more than half a point, imply that this effect can be of major consequence for a large share of tactical voters, potentially changing the order of their preferences over time. The effect might be especially important if it cumulates over several elections, but, in this study, the effect is identified on a short-term basis, and it cannot be concluded that it also persists in the long-term. That remains a topic for further research.

The results are consistent with the predictions of cognitive dissonance theory, which at least for the voters studied here suggests that dissonance does play a role in producing the effect of voting on party preferences. These findings may therefore call for a revision of the traditional interpretation of the role of voting based on the notion of habit formation. In contrast to the traditional understanding in political science, cognitive dissonance has the interesting implication that voters who are led to vote for parties they do not prefer may still come to like them. Thus, parties that gain from tactical voting due to favorable electoral rules may in the long-run also gain more genuine support as voters come to identify with the parties they vote for.

This mechanism may have important implications for the long-run development of party systems, disfavoring new and small parties—even when they offer policies many voters would otherwise like. A possible illustration from the UK is the electoral district of Boston and Skegnee, where the Conservative candidate narrowly defeated the Labour Party both in the 1997 and 2001 elections. Interestingly, the Liberal Democrats persistently lost popular support during this time: from 16.6 % in 1997 to 12.4 % in 2001, and 8.7 % in 2005. It is likely that the extremely competitive environment accounts at least in part the electoral fortune of the Liberal candidates, by causing tactical voting, but we can neither rule out that former Liberal supporters who strategically defected came to like the party they voted for, increasing their likelihood of voting for the same party again—and such attitudinal change may also have contributed to the declining Liberal vote over these elections.

The effect we identify may similarly have important implications for the consolidation of new party systems. Some institutional constraints, such as the imposition of high electoral thresholds or the adoption of mixed-member systems, generate incentives to cast tactical votes. Such votes may in turn alter people's party preferences. Support for small parties may gradually wear off to the advantage of larger parties that are more likely to gain political influence. Various post-Communist countries in Eastern Europe, such as Romania and Poland, who witnessed a plethora of parties and considerable electoral volatility in their first

democratic elections, are telling examples. In these countries, more stable party support began to evolve, at least in part as a result of the introduction of incentives to refrain from voting for small parties (see, for example, Tavits 2005). Our results suggest that such induced tactial voting over time may generate genuine party ties and promote the creation of stable party systems.

Acknowledgment The authors are grateful for comments by three anonymous reviewers, the editors of this journal, Juan A. Mayoral, and the participants of the Elections, Public Opinion and Parties (EPOP) Annual Conference at the University of Essex, September 10–12, 2010. The usual disclaimer applies.

References

Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of Econometrics, 113*(2), 231–263.

Abadie, A., & Imbens, G. (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica, 74*(1), 235–267.

Alvarez, R. M., & Nagler, J. (2000). A new approach for modeling strategic voting in multiparty elections. *British Journal of Political Science, 30*, 57–75.

Anderson, C. J., Mendes, S. M., & Tverdova, Y. V. (2004). Endogenous economic voting: Evidence from the 1997 British Election. *Electoral Studies, 23*(4), 683–708.

Anderson, C. J. (2007). The end of economic voting? Contingency dilemmas and the limits of democratic accountability. *Annual Review of Political Science, 10*, 271–296.

Angrist, J. D., Imbens, G., & Rubin, D. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association, 91*(434), 444–472.

Aronson, E., Fried, C., & Stone, J. (1991). Overcoming denial and increasing the intention to use condoms through the induction of hypocrisy. *American Journal of Public Health, 81*(12), 1636–1638.

Arulampalam, W., Booth, A. L., & Taylor, M. P. (2000). Unemployment persistence. *Oxford Economic Papers, 52*(1), 24–50.

Bem, D. (1972). Self-Perception Theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 1–62). New York: Academic Press.

Besley, R., & Joslyn, M. (2001). Cognitive dissonance and post-decision attitude change in six presidential elections. *Political Psychology, 22*(3), 521–540.

Brader, T., & Tucker, J. A. (2001). The emergence of mass partisanship in Russia, 1993–1996. *American Journal of Political Science, 45*(1), 69–83.

Brehm, J. W. (1956). Post-decision changes in desirability of alternatives. *Journal of Abnormal and Social Psychology, 52*(3), 384–389.

Brody, R. A. (1991). Stability and change in party identification: Presidential to off-years. In P. M. Snidman, R. Brody, & P. E. Tetlock & (Eds.), *Reasoning and choice: Explorations in political psychology* (pp. 179–205). Cambridge: Cambridge University Press.

Butler, D., & Stokes, D. E. (1974). *Political change in Britain* (2nd edn.). London: MacMillan.

Cain, B. (1978). Strategic voting in Britain. *American Journal of Political Science, 22*(3), 639–655.

Campbell, A., Converse, P. E., Miller, W. E., & Stokes D. E. (1960). *The American Voter.* New York: John Wiley & Sons, Inc.

Converse, P. E. (1969). Of time and partisan stability. *Comparative Political Studies, 2*(2), 139–171.

Converse, P. E. (1976). *The dynamics of party support: Cohort-analyzing party identification.* Beverly Hills: Sage Publications.

Cowden, J. A., & McDermott, R. M. (2000). Short-term forces and partisanship. *Political Behavior, 22*(3), 197–222.

Curtice, J., & Steed, M. (1988). Appendix 2: analysis. In D. Butler & D. Kavanaugh (Eds.), *The British General Election of 1987* (pp. 316–362). New York: St. Martin's Press.

Daniel, S. R., Armstrong, K., Silber, J. H., & Rosenbaum, P. R. (2008). An algorithm for optimal tapered matching, with application to disparities in survival. *Journal of Computational and Graphical Statistics, 17*(4), 914–924.

Diamond, A., & Sekhon, J. S. (2008). *Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies, working paper*. Berkeley: University of California.

Dobson, D., & St Angelo, D. (1975). Party identification and the floating vote: Some dynamics. *American Political Science Review, 69*(2), 481–490.

Egan, L. C., Santos, L. R., & Bloom, P. (2007). The origins of cognitive dissonance: Evidence from children and monkeys. *Psychological Science, 18*(11), 978–983.

Eggers, A. C., & Hainmueller, J. (2009). MPs for sale? Returns to office in postwar British politics. *American Political Science Review, 103*(4), 513–533.

Eiser, R. J. (1992). Attitudinal judgement: The to and fro of assimilation-contrast. In D. Granberg & G. Sarup (Eds.), *Social judgement and intergroup relations. Essays in honor of Muzafer Sherif* (pp. 129–146). New York: Springer.

Evans, G. A. (1994). Tactical voting and Labour's prospects. In A. Heath, R Jowell, & J. Curtice & (Eds.), *Labour's last chance?: The 1992 Election and beyond* (pp. 65–84). Aldershot: Dartmouth.

Evans, G. A., & Andersen, R. (2006). The political conditioning of economic perceptions. *Journal of Politics, 68*(1), 194–207.

Evans, G. A., Curtice, J., & Norris, P. (1998). New Labour, new tactical voting? The causes and consequences of tactical voting in the 1997 General Election. In D. Denver, J. Fisher, P. Cowley, C. Pattie (Eds.), *British Elections and Parties Review vol. 8: The 1997 General Election* (pp. 65–79). London: Frank Cass.

Evans, G. A., & Heath, A. F. (1993). A tactical error in the analysis of tactical voting: A response to Niemi, Whitten and Franklin. *British Journal of Political Science, 23*(1), 131–137.

Festinger, L., Riecken, H. W., & Scachter, S. (1956). *When prophecy fails*. Minneapolis: University of Minnesota Press.

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology, 58*, 203–210.

Festinger, L. (1957). *A Theory of cognitive dissonance*. Evanston: Row, Peterson.

Fiorina, M. P. (1981). *Retrospective voting in American Elections*. New Haven: Yale University Press.

Fisher, S. D. (2004). Definition and measurement of tactical voting: The role of rational choice. *British Journal of Political Science, 34*, 152–166.

Fisher, S. D., & Curtice, J. (2006). Tactical unwind? Changes in party preference structure and tactical voting in Britain between 2001 and 2005. *Journal of Elections, Public Opinion and Parties, 16*(1), 55–76.

Franklin, C. (1984). Issue preferences, socialization and the evolution of party identification. *American Journal of Political Science, 28*(3), 478–479.

Franklin, C., & Jackson, J. E. (1983). The Dynamics of party identification. *American Political Science Review, 77*(4), 957–973.

Gerber, A. S., & Green, D. P. (2000). The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment. *American Political Science Review, 94*(3), 653–663.

Gerber, A. S., & Green, D. P. (2005). Correction to Gerber and Green (2000), replication of disputed findings, and reply to Imai (2005). *American Political Science Review, 99*(2), 301–313.

Green, D. P., & Palmquist, B. (1994). How stable is party identification?. *Political Behavior, 16*(4), 437–466.

Green, D. P., & Palmquist, B. (1990). Of artifacts and partisan instability. *American Journal of Political Science, 34*(3), 872–902.

Grimmer, J., Hersh, E., Feinstein, B., & Carpenter, D. (2011). *Are close elections random*? Unpublished Manuscript. http://www.stanford.edu/~jgrimmer/CEF.pdf.

Groenendyk, E. (forthcoming). Justifying party identification: A case of identifying with the "lesser of two evils". *Political Behavior*.

Grofman, B., Wayman, F., & Barreto, M. (2009). Rethinking partisanship: Some thoughts on a unified theory. In J. Bartle & P. Bellucci (Eds.), *Political parties and partisanship: Social identity and individual attitudes* (pp. 60–73). Routledge: Ruhr University Bochum.

Hansen, B. B., & Bowers, J. (2008). Covariate balance in simple, stratified and clustered comparative studies. *Statistical Science, 23*(2), 219.

Hansen, B. B., & Bowers, J. (2009). Attributing effects to a cluster-randomized get-out-the-vote campaign. *Journal of the American Statistical Association, 104*(487), 873–885.

Heath, A., Jowell, R., Curtice, J., Evans, G. A., Field, J., & Witherspoon, S. (1991). *Understanding political change: The British voter 1964–1987*. Oxford: Pergamon Press.

Heider, F. (1958). *The psychology of interpersonal relations*. New York: John Wiley Sons.
Howell, S. (1980). The behavioral component of changing partisanship. *American Politics Quarterly, 8*(3), 279–302.
Howell, S. (1981). Short term forces and changing partisanship. *Political Behavior, 3*(2), 163–180.
Huber, G., Gerber, A., & Sempolinski, J. (2009, 3–6 September). *Using ballot order to test for cognitive dissonance: Results from a natural experiment*. Paper presented at the 2009 APSA Annual Meeting, Toronto.
Imai, K. (2005). Do get-out-the-vote calls reduce turnout? The importance of statistical methods for field experiments. *American Political Science Review, 99*(2), 283–300.
Jackson, J. E. (1975). Issues, party choices, and presidential votes. *American Journal of Political Science, 19*(2), 161–185.
Johnston, R. (1981). Regional variations in British voting trends—1966-1979: Tests of an ecological model. *Regional Studies, 15*(1), 23–32.
Johnston, R., & Pattie, C. (1991). Tactical voting in Britain in 1983 and 1987: An alternative approach. *British Journal of Political Science, 21*, 95–128.
Kern, H. L., Hainmueller, J. (2009). Opium for the masses: How foreign media can stabilize authoritarian regimes. *Political Analysis, 17*(4), 377–399.
Kiewiet, R. (2009, 28–30 August). *The ecology of tactical voting in Britain, 1983–2005.* Paper presented at the 2009 EPOP conference, University of Strathclyde, Glasgow.
Knoke, D. (1976). *Change and continuity in American politics: The social basis of political parties*. Baltimore: Johns Hopkins University Press.
Larsen, R. J., & Buss, D. M. (2005). *Personality psychology: Domains of knowledge about human nature*. Boston: McGraw Hill.
Lu, B., & Rosenbaum, P. R. (2004). Optimal pair matching with two control groups. *Journal of Computational and Graphical Statistics, 13*(2), 422–434.
Markus, G. B., & Converse, P. E. (1979). A dynamic simultaneous equation model of electoral choice. *American Political Science Review, 73*(4), 1055–1070.
McPhee, W., & Ferguson, J. (1962). Political immunization. In W. N. McPhee & W. A. Glaser (Eds.), *Public opinion and congressional elections* (pp. 155–179). Glencoe: The Free Press.
Meredith, M. (2009). Persistence in political participation. *Quarterly Journal of Political Science, 4*(3), 186–208.
Mullainathan, S., & Washington, E. (2009). Sticking with your vote: Cognitive dissonance and political attitudes. *Applied Economics, 1*(1), 86–111.
Niemi, R., Whitten, G., & Franklin, M. N. (1993). People who live in glass houses. *A response to Evans and Heath's critique of our note on tactical voting. British Journal of Political Science, 23*, 549–563.
Plumb, E. (1986). Validation of voter recall: Time of electoral decision making. *Political Behavior, 8*(4), 301–312.
Robinson, G., McNulty, & J. E., Krasno, J. S. (2009). Observing the counterfactual? The search for political experiments in nature. *Political Analysis, 17*(4), 341–357.
Rosenbaum, P. R. (2010). *Design of observational studies*. New York: Springer.
Sekhon, J. S., & Mebane, W. R. Jr. (1998). Genetic optimization using derivatives: Theory and application to nonlinear models. *Political Analysis, 7*(1), 189–203.
Sherif, M., & Hovland, C. I. (1961). *Social judgement: Assimilation and contrast effects in communication and attitude change*. New Haven: Yale University Press.
Shively, W. P. (1970). The elusive 'psychological' factor: A test for the impact of electoral systems on voters' behavior. *Comparative Politics, 3*(1), 115–125.
Spafford, D. (1972). Electoral systems and voters' behavior: Comment and a further test. *Comparative Politics, 5*(1), 129–134.
Stone, J. R. (2000). Introduction. In J. R. Stone (Eds.), *Expecting armageddon: Essential readings in failed prophecy* (pp. 1–30). New York: Routledge.
Tavits, M. (2005). The development of stable party support: Electoral dynamics in post-communist Europe. *American Journal of Political Science, 49*(2), 283–298.
Van der Brug, W. (2001). Perceptions, opinions and party Preferences in the face of a real world event: Chernobyl as a natural experiment in political psychology. *Journal of Theoretical Politics, 13*(1), 53–80.
Zanna, M. P., & Cooper, J. (1974). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. *Journal of Personality and Social Psychology, 29*(5), 703–709.