Does campaign contact mobilise individual voters? A quasi-experimental approach

Charles Pattie
Department of Geography
University of Sheffield

Adam Whitworth
Department of Geography
University of Sheffield

Ron Johnston
School of Geographical Sciences
University of Bristol

Does campaign contact mobilise individual voters?
A quasi-experimental approach

Abstract

Analyses of local campaign effects are dominated by aggregate-level analyses of constituency activity. Though individual-level data are available on whether voters are (or remember being) contacted by parties during campaigns, their analysis is fraught with difficulties, not least the extent to which memory of campaign contact is itself conditioned partly on party allegiance, creating a circularity in the analysis of the impact of party contact on vote choice. To some degree, this can be (and has been) dealt with in a regression framework. However, this does not fully deal with the potential difficulties. Ideally, more experimental approaches are needed to tease out definitively the effects of campaign exposure on individual's election decisions. However, that is fraught with some difficulties. In this paper, therefore, we examine alternative estimation procedures which provide quasi-experimental estimates of campaign effects.

The now-voluminous literature on constituency campaign effects in Great Britain is dominated by aggregate, constituency-level analyses. Most studies look at how changes in parties' constituency vote shares correlate with measures of aggregate constituency campaign effort (whether measured by campaign spending data, information supplied by party election agents, surveys of party members, or some permutation of these: for a recent example, see Johnston *et al*., 2011). As a result of this research, it is now widely accepted that local campaigns pay electoral dividends for parties. However, while the aggregate relationship is well-established, demonstrating the micro-foundations of the constituency campaign is more problematic: how might we uncover the impact of the local campaign on individual voters?

Most attempts to do so utilise data from election surveys which ask respondents whether they were contacted by individual parties' constituency campaigns during the election. Those reporting being contacted by a party are, other things being equal, more likely to vote for that party than are those who do not report such contacts (e.g. Denver et al., 2004; Pattie and Johnston, 2010; Clarke *et al.*, 2004, 2009; Johnston, Cutts et al., 2012). What is more, self-reported measures of campaign contact correlate well with other measures of campaign intensity, suggesting they are good indicators of exposure to the campaign: voters are more likely to report being contacted by a party during an election campaign if they live in a constituency where that party mounted an active campaign than where its campaign was less intense (Pattie *et al*., 1994; Denver *et al*., 2004; Johnston, Pattie *et al*., 2012).

Using self-reported campaign exposure raises serious methodological problems, however, as respondents' recollections of being contacted during a campaign are not independent of their partisan leanings. On the contrary, individuals are more likely to remember being contacted by a party they already support than by one they do not favour. This introduces a threat of substantial selection biases: any apparent correlation between self-reported contact from a campaign and vote may simply be an artefact of the tendency for those already pre-disposed to a party (and hence likely to vote for it anyway) being the most likely to remember being contacted by it.

Although there is very strong and extensive circumstantial evidence of substantial campaign effects in Great Britain, therefore, there is a need for more robust micro-level foundations. One strategy which would get around the selection bias difficulties inherent in conventional observational survey research would be to employ large-scale field experiments, similar to those now widely used to analyse non-partisan get-out-the-vote measures (see e.g. Green and Gerber, 2004). Unfortunately, such experiments are more difficult to conduct where there are partisan implications. But it is possible, with careful analysis of well-constructed individual survey instruments, to move beyond aggregate studies and to attack the problem through quasi-experimental approaches. In this paper, we examine the potential of two types of quasi-experimental method designed to reveal the impact of policy interventions on key outcomes.

A quasi-experimental approach

A key problem in using survey data to analyse the impact of campaign activity is how to deal with the endogeneity of self-reported party campaign contact and vote intention. The potential scale of the issue can be gauged by looking at data from the 2010 British Election Study Campaign Internet Panel, which interviewed a large national sample before, during and after the election campaign (13,334 individuals answered both the pre- and post-election

waves of the survey).[1] Respondents to the post-election wave were asked how they had voted in the election, and whether they had been contacted by the parties during the campaign. Those reporting being contacted were asked which party (ies) had contacted them, and how the contact had been made: the alternatives on offer included: being telephoned; having leaflets or other forms of mail delivered to their home; being visited at home (most such contacts will have been from a party canvasser); being contacted in the street; receiving a message via Twitter or some other online social networking system; receiving a text message; or via some other means. (Unfortunately, the survey did not discern how many times each respondent had been contacted through each medium.) Not surprisingly, relatively large numbers of respondents reported receiving leaflets from the major parties: 38% from Labour, 42% from the Conservatives, and 38% from the Liberal Democrats. Very few indeed reported being contacted via new social media.

Our main interest here, however, is in those who reported some form of face-to-face contact with a party's campaign, whether at home or in the street (7% of respondents reported being visited at home by Labour and 2% reported being contacted by the party in the street; the equivalent figures for contact by the Conservatives were 9% and 3% respectively; for the Liberal Democrats, they were 5% and 2%). Canvassing voters in their homes is amenable to more careful targeting than street contacts. While the former are directed at known constituency residents – or even residents of target neighbourhoods in the constituency based, increasingly, on geodemographic profiling of residents' characteristics and expenditure habits (Farrell and Webb, 2002; Cutts, 2006; Fisher and Denver, 2008) – the latter, even though they are likely to take place in the same competitive seats, cannot, by their nature, guarantee that those contacted will be voters in the 'correct' constituency. Many street contacts, for instance, will be generated by activists stopping individuals in shopping areas and some at least of those approached in such areas, especially within large cities, will reside outside the local constituency. There is likely to be a greater degree of randomness, therefore, to street contacts than to contacts with voters at home. That said, both forms of contacting voters are liable to be more common in competitive than in uncompetitive seats. We therefore concentrate on these face-to-face contacts as they could by and large be attributed to local

[1] The BES Campaign Internet Panel Survey (CIPS) interviewed the same group of individuals shortly before and immediately after the 2010 UK General Election, allowing us to assess the impact of events during the campaign period, such as being contacted by the party campaigns, on vote choice. In addition, the CIPS data offers several advantages over conventional face-to-face surveys. The sample size is much larger than most face-to-face surveys: 16,816 individuals were interviewed in the pre-election wave of the survey, and 13,356 were interviewed in the post-election wave. Furthermore, the pre- and post-election waves were in the field for only a few days. The first responses to the pre-election wave came in on 29 March 2010, and the last on 7 April 2010: all replies to the post-election wave were captured between 7 and 24 May 2010 (and 90% had been returned by 14 May, barely a week after polling day). By contrast, large face-to-face surveys can take some months to complete (the BES 2010 face-to-face survey's pre-election wave began interviewing on 23 January 2010 and continued till 19 April 2010, while interviews for the post-election wave took place between 9 May and 16 August, over 3 months after the election). The longer a survey is in the field, the greater the risk that events during the fieldwork will contaminate respondents' answers. In 2010, the election outcome was inconclusive and five days passed before a coalition could be formed between the Conservatives and the Liberal Democrats. That coalition proved highly controversial, especially among Labour and Liberal Democrat voters (few of whom had anticipated that the Liberal Democrats would form a government with the Conservatives). There is a distinct possibility, therefore, that responses after the formation of the coalition might be systematically biased by that event. The quicker the survey responses were gathered, therefore, and the more that were obtained between the election and the emergence of the coalition, the better. Clearly, the CIPS data has a major advantage in this respect.

campaign efforts in the constituencies.[2] This was not always true of contacts via electronic media given that some of these were co-ordinated by the party nationally with no necessary linkage to the constituency campaign. And while considerable effort was expended on telephone campaigns to complement local canvassing, for instance, some telephone calls will have been to party supporters, eliciting financial support, rather than to voters in key battleground constituencies. In any case, contact via telephones and new social media were relatively rare. Leafleting, meanwhile, remains a largely local, quasi-random activity – leaflets are distributed to all homes in an area (under the 1983 *Representation of the People Act*, all candidates are entitled to mail one election address free of charge to every elector during a campaign – though the parties must cover the printing costs – and parties, if they have the resources to do so, can send out more material over and above this). But the relative ubiquity of campaign leaflets and the undoubted tendency for most to move direct from the doormat to the rubbish bin with only the most limited of scrutiny by the voter renders exposure to them a rather crude indicator. In any case, as discussed below, the analyses later in the paper control for party constituency spending. As printing costs accounted for 85% of local party expenditure during the 2010 official 'short' campaign (and the bulk of that expenditure was on leaflets and similar material), this largely captures the effect of leafleting.[3]

Using the BES we can classify respondents according to how they voted and by whether or not they were contacted face to face by each of the major parties. The results would seem to imply substantial campaign effects (table 1). For instance, 53% of those who reported being contacted face-to-face by the Conservatives also reported voting for the party compared to only 31% of those who did not report being contacted, suggesting a substantial effect size of 22 percentage points. Similarly impressive-looking effects are evident for the other parties too: Labour's vote share was 53% among those who remembered an encounter with the party, compared to 21% among those who did not (an effect size of 32 percentage points), while the comparable figures for the Liberal Democrats were 50% and 23% (for an effect size of 27 percentage points).

Not surprisingly, as evidence of local campaign effects, this is too good to be true. Several factors contribute to deceptively impressive-looking effect sizes. To some extent, selective memory operates: people are more likely to remember contact by a party they favour than by one they oppose or are indifferent to. To some extent, too, voting for a party will be much more common among those reporting being contacted by it than among those who do not report this because of the actions of the parties themselves. One of the purposes of the constituency campaign is to mobilise supporters rather than to convert opponents. Parties do not waste much time chasing voters they already suspect are unlikely ever to vote for them (as this is liable only to boost their rivals' votes). Rather, they try to ensure that as many of their own supporters turn out as possible. So if parties are more likely to contact those already 'in' their camp to ensure they turn out, it is hardly surprising to find that contactees vote for

[2] The analyses reported below utilise a measure of face-to-face campaigning which combines both contact at home and contact in the street. We have also repeated all analyses utilising only reported contact at home. The results, not reproduced here, replicate the findings reported in this paper.

[3] We have, however, repeated our analyses with the measures of self-reported campaign contact by each party extended to include not only contacts at home or in the street, but also via leaflets. The results, not reported here, are entirely consistent with our overall argument. Effect sizes for the more extensive measure of self-reported contact are generally smaller than for the face-to-face contact measures. But all are significant and in the expected directions.

the party at much higher rates tan those not contacted (even if there were no selective memory effects).

Both effects will therefore tend to exaggerate the apparent effectiveness of campaign contact. There is some evidence of this in the survey data; those already leaning towards a party before the official election started were rather more likely to recall being contacted by that party during the campaign than were those who were not pre-disposed towards it. This is illustrated in table 2 which displays the relationship between what individuals said about their partisan affiliations in the pre-election wave of the survey with whether, in the post-election wave, they reported being contacted by each of the three main parties during the campaign. In every case, those reporting being contacted by a party were much more likely to already be supporters of the party before the contest than were those who did not report being contacted. For instance, 46% of those who said they had been contacted by the Conservatives were already Conservative identifiers, compared to only 28% of those who did not remember a contact. Similar disparities exist for contact by the other parties. The gap in identification with the Liberal Democrats between those remembering a Liberal Democrat contact and those not recalling one is narrower than the equivalent gaps for the other parties, but this is largely because there were just fewer Liberal Democrat identifiers to begin with: this notwithstanding, the gap is in the same direction as for the other parties. Whether this is a result of selective memory, of campaign targeting, or possibly both, the implication is that a simple examination of the difference in the percentage voting for a party between those who report being contacted by it and those who do not is very likely to substantially over-estimate the true size of any campaign effect at the level of individual voters.

In addition, at least some respondents are likely to report campaign contacts which never occurred or votes they did not cast. In some cases, this will be a consequence of simple mistaken memory. In other cases, it may be due to the tendency for people to give answers they feel are socially desirable: rather than appear to be inattentive and disengaged from the campaign, for instance, some are liable to mis-report the degree of attention they did pay to it. And this will also introduce some bias and inaccuracy into survey-based estimates of campaign effects.

So how might we deal with this? The gold standard is to employ randomized field experiments, which have been employed with great success to look at the impact of get-out-the-vote campaigns (Green and Gerber, 2004; John and Brannan, 2008). Typically in such studies, participants are assigned at random to different groups: one group might receive no messages encouraging them to vote; another might receive a message emphasising the importance of voting, delivered face-to-face by a get-out-the-vote campaign; yet another group might receive a telephone call from the get-out-the-vote campaign. By using post-election official records of who did and who did not vote, it then becomes possible to ascertain whether different forms of campaign contact are more or less likely to persuade individuals to vote than no contact at all. And as individuals are randomly assigned to groups, differences in turnout between groups can reasonably be attributed to the treatments themselves (i.e. to different modes of campaigning) rather than to the sorts of individuals in each group. Many such studies have now been conducted, almost all confirming the importance of get-out-the-vote messages (particularly when delivered in person, and especially when reinforced with information which stresses that participation is a local norm: Gerber *et al*., 2008, 2010; Davenport *et al*., 2010; Panagopoulos, 2010).

Comparisons of field experiments and conventional surveys employing respondents' self-reported exposure to campaigns and votes suggest that the latter very substantially overestimate the effectiveness of campaigns. Vavrek (2007) provides some insights. Her study was based on a randomized field experiment involving 1500 respondents, who were randomly assigned either to a group which did, or to one which did not receive a get-out-the-vote message (hence giving an unbiased measure of who was exposed to the campaign): participants' individual turnout was also verified against official records. At the same time, a conventional self-report survey was also administered to the same individuals, asking them for their self-reported accounts of whether they had been exposed to the campaign or had voted. Her results suggest that estimates based on survey self-reports exaggerate effect sizes very substantially (p. 326): a 1% effect size in the randomized field experiment becomes 7% in the self-reported survey.

In principle, therefore, one might consider a similar controlled field experiment to investigate the impact of different partisan messages on vote choice in which individuals are randomly assigned to different groups and with the nature of the campaign contact varying from group to group: a control group might be left untouched, with no contact from any party during the campaign; one group might be canvassed by Labour alone; another by the Conservatives, and so on. Indeed, some early studies of campaign effects came close to this (an approach adopted in pioneering work by Bochel and Denver, 1971). ~~For instance, Bochel and Denver (1971) conducted an experiment during local council elections in Dundee. Two tower blocks in a safe Labour-controlled ward were selected for the experiment: the blocks were broadly similar in terms of their social composition. Residents in one block were subjected to intensive canvassing on behalf of the Labour party; residents in the other block of flats were not canvassed. After the election, a survey of residents in both towers revealed that both turnout and the Labour vote share were higher in the canvassed than the uncanvassed block.~~ However, ~~unlike Gerber and Green's get-out-the-vote experiments, assignment to treatment groups was not truly random. Furthermore, while Gerber and Green were able to confirm~~ actual turnout can be confirmed, it is not possible to ascertain from official election returns which parties particular individuals voted for: ~~Bochel and Denver's~~ experimental designs ~~(and all similar experiments)~~ therefore have~~d~~ to rely on self-reported voting in a post-election survey. A threat of selection bias therefore remains. Furthermore, ~~interesting and suggestive though it was, Bochel and Denver's experiment was inevitably restricted to just one location and to a low-intensity local election in a safe Labour ward where the experiment would have no impact on the outcome, making it difficult to be sure the findings would be replicable on a wider stage. T~~the prospects of obtaining agreement from the political parties to ~~extend~~ conduct such a study ~~to~~ in ~~the wider context of~~ a general election (which would require their agreement to give up control of their own campaigns in at least some key battleground seats) are, it must be said, limited in the extreme.

An alternative approach is to conduct a lab experiment, in which the entire research process takes place in a controlled environment (e.g. Ansolabehere and Iyengar, 1995; Norris *et al.*, 1999). ~~Ansolabehere and Iyengar's (1995) well-known study of the effectiveness of negative political advertising adopts just such a strategy. A sample of 3,000 California voters was employed for their experiment. Individuals were placed at random in different treatment groups, each of which was asked to watch a television news broadcast, within which there were commercial breaks. A control group was shown a version of the broadcast in which no political ads ran during the commercial break. Other treatment groups were shown some political commercials embedded within the news broadcast. The tone (positive or negative) of the commercials was varied from group to group. Participants were also interviewed before~~

~~and after they viewed the commercials, allowing the research team to assess how much, and in which direction, the commercials altered peoples' opinions.~~ This does allow for genuinely random assignment of individuals to treatment groups. However, because they are conducted under laboratory conditions rather than in the midst of real elections, such studies suffer from inevitable doubts regarding their wider applicability~~:. After all,~~ no matter how carefully constructed the experiment is it cannot replicate an actual vote decision as nothing depends on the participants' opinions. There is an obvious concern, therefore, over the face validity of their results.

So how can we move forward? Interestingly, Vavrek's (2007) study, while showing that self-reports can lead to substantial over-estimates of campaign effects, also suggests a way out: regression models predicting self-reported turnout from self-reported campaign contact, but also including control variables for factors such as interest in politics (which might be related to the tendency to over-report both voting and exposure to the campaign), the effect sizes become comparable to those produced in the randomised experiment. In this paper, we compare three different strategies for the analysis of individual-level campaign effects using conventional survey data.

All three provide estimates of the impact of an intervention on behaviour while trying to deal with potential biases caused by variations between treatment and control groups, including possible selection biases. For instance, we might be interested in how much support for party X changes over the course of an election campaign, and in particular in whether being contacted by that party stimulates support. Ideally, we would follow two groups of voters throughout the campaign, one (the treatment group) which is contacted by the party and the other (the control) which is not, and we measure intention to vote for the party before and after the election for both groups. For the sake of argument, let us say that support for party X among the treatment group rose from 30% to 40% over the course of the campaign, a 10 percentage point increase, while among the control group support rose from 25% to 32%, a 7 percentage point rise (figure 1). A simple approach to assessing the effect of contact would compare the start- and end-points for the treatment group (the vertical distance between points A and B in figure 1). So exposure to the campaign raised support for X by 10 percentage points. But there is no control here. Perhaps it would be better to compare the post-election levels of support for party X in the treatment group with the same for the control group (the vertical distance between D and B in figure 1)? In this example, this yields a more modest estimate of the campaign effect: exposure to the campaign raised party X's vote by 8 percentage points. But a moment's reflection reveals the flaw: support for party X rose in *both* the treatment (the rise from A to D) *and* the control group (the rise from C to D) over the course of the campaign, and the treatment group was already relatively more predisposed towards party X than the control before the campaign began. So some of the increase in support for X among the treatment group might have happened anyway, even if there had been no campaign contact. We need to partial that out before we can get closer to isolating the effect of the campaign. So support for X in the control group went up by 7 percentage points over the course of the campaign. Even if the treatment group was not contacted, therefore, it is quite possible that it, too, would have seen support for X go up by the same amount, rising from 30% to 37% (represented in figure 1 by the dotted line from A to E). The gap between the latter figure and the actual post-election level of support for X (40%) gives us our best estimate of the independent effect of the campaign (the gap between

E and B in figure 1, a 3 percentage point rise), as it partials out the general drift in support for the party.[4]

Two of the three strategies examined below employ difference-in-difference (DiD) approaches. The first, and most widely applied, of these is to use a conventional regression model to hold constant other relevant influences on vote choice. As this is the default option in many existing studies there is little novel here (see e.g. Clarke *et al.*, 2004, 2009; Pattie and Johnston, 2010; Fisher *et al.*, 2011). But it does provide a benchmark against which to compare the results of our other methods: below, we discuss this as the 'conventional regression' approach. However, despite its simplicity and ubiquity, the potential for selection biases, discussed above, presents a serious problem for this widely-used approach.

The second of the two difference-in-difference estimation methods used here exploits the fact that each respondent to the 2010 BES CIPS was interviewed both before and after the election by including both observations as separate cases in the data set and employs an adapted regression modelling approach in which the key variable of interest is an interaction between the time period (before or after) and exposure to the intervention, in our case campaign contact (Machin et al., 2004; Meghir and Palme, 2005; Wilkinson and McLennan, 2012; Ikenwilo, 2013). For reasons explained below, we refer to this as the 'DiD with interaction' approach.

The final method examined here, propensity score matching (PSM), uses a somewhat different strategy. It builds and uses a propensity score to match individuals exposed to the intervention to individuals not exposed to the intervention and then assesses differences in outcomes between the two groups (Dehejia and Wahba, 2002; Bryson et al., 2002; Jalan and Ravallion, 2003). In the absence of data from genuine experimental work both approaches have the potential to provide purer and more reliable estimates of campaign effect than the simple 'conventional regression' approach to statistical control.

All of the methods compared here use regression models at some stage, whether to produce the main analysis (the conventional regression method, and DiD with interaction) or as a step towards producing matched samples for comparison (PSM). To ensure comparability, therefore, the same explanatory variables are used throughout. So what might affect whether an individual reports being contacted face-to-face by a party during the election campaign?

In part, this will be driven by pre-existing predispositions towards a party. Parties are liable to concentrate their attention on those who already support them, or who are undecided (for instance, floating voters, previous abstainers, or new voters) rather than on those who are likely to vote for their rivals (why mobilise those who will only vote against you?). Through canvassing activity in the months and years before an election they will have some idea of who many of these individuals are, or at least the type of neighbourhood in which most of them live (Johnston, Cutts et al., 2012). In addition, there is a strong likelihood that individual voters will be more likely to recall and report being contacted by a party they already support

[4] Formally, we can write the difference-in-difference estimate as

$$\text{Impact} = (Y_{t1} - Y_{t0}) - (Y_{c1} - Y_{c0})$$

where Y_{i0} is pre-election probability of voting for a party, Y_{i1} is the probability of actually voting for it, t is the treatment group (those reporting being contacted by the party during the campaign) and c is the control group (those who report no contact).

than being contacted by one they do not. For both reasons, we include independent variables to take into account self-reported vote at the previous General Election in 2005 and voters' party identification on the eve of the 2010 election. As a consequence, we do not need to include further explanatory variables for the 'standard' socio-demographic and --economic measures often included in models of individual behaviour – such as class, age, gender and so on – as they are themselves related to party identification and past voting: the latter variables in effect take into account the former. A more parsimonious model without these additional independent variables therefore misses little. In addition, parties concentrate their constituency campaign efforts in those seats in which they face the most intense competition: they put only token effort into seats where they are bound to lose and relatively little effort into those they are bound to win (Pattie and Johnston, 2003; Johnston *et al.*, 2013). Hence we expect individuals to be more likely to report being contacted by a party in seats where that party is campaigning hard than in seats where it is making less effort. To capture this we include independent variables for the amount each party spent on its 2010 'short campaign' (the period from the dissolution of Parliament to the date of the election) in each respondent's constituency as a percentage of the legal maximum expenditure permitted there. It is also likely that those who actively contact their political representatives will be more likely to pay attention to the campaign nationally through the mass media and to be included in parties' data bases of voters to contact during the weeks immediately preceding the election compared with voters who do not contact politicians. We capture this using a question in the pre-election wave of the BES asking respondents whether they had sought help from their local MP. Finally, we expect that parties will be more likely to contact (and the contact is more likely to be remembered by) those who pay close attention to politics than those who pay little or no attention. To capture this we add responses to a question from the pre-election wave of the survey, asking individuals to rate, on an 11-point scale, how much attention they pay to politics (the responses are coded so that high scores indicate most attention).

So far, so conventional: difference-in-difference estimates using the conventional regression approach

As described above, our first cut at reaching a more nuanced estimate of the electoral impact of face-to-face contact on vote uses a conventional regression approach to take into account other possible influences on vote choice (table 3: in these and all subsequent regression models, we report robust standard errors). Although the dependent variable is binary in form, we employ OLS regression here. This is because the goal for this analysis is to provide a benchmark against which to compare the DiD with interaction and PSM results introduced later in the paper and in which the key 'impact' estimator of interest is an interaction term of time and treatment. Given that interaction terms in logit or probit models are not interpretable in the same way as in linear models (Ai and Norton, 2003), and hence would not capture the true difference-in-difference with interaction estimate, the application of linear difference-in-difference models to binary outcomes is commonplace (Weinick et al., 2000; Fu et al., 2007; Liu et al, 2010). For comparability, therefore, OLS is employed at this stage too, although we have checked our results using logit models and the key findings are the same in both specifications.

The results are straightforward and largely as expected. Not surprisingly, those already predisposed towards a party (whether because they identified with it or had voted for it at the previous election) were generally more likely to vote for it than were those who were not predisposed towards it. And the more actively a party campaigned in a respondent's constituency (as indexed by the party spending variables) the more likely that person was to

vote for the party. In addition, there is some evidence that parties' campaign efforts also discouraged people from voting for their rivals. The harder the Conservatives worked in a constituency, for instance, the less likely it was that respondents living there would vote Labour. Similarly, the bigger the effort put in by the Liberal Democrat candidate the less likely it was that individuals would vote either Labour or Conservative, and the harder the Labour candidate worked the less likely respondents were to report voting Liberal Democrat.

Our primary focus, however, is on the coefficients for face-to-face contact. In every case, reporting being contacted by a party has a significant and positive effect on the likelihood of voting for that party, even when we take into account other influences on the vote decision. In other words, the conventional analysis upholds the conventional wisdom: being contacted by a campaign encourages individuals to vote for the relevant party.

There are some interesting differences between the parties, however: the Labour face-to-face contact coefficient is almost twice as large as that for the Conservatives, while the coefficient for reporting being contacted by the Liberal Democrats is between the coefficients for the other two parties (though closer in size to that for the Conservatives than for Labour). Since the dependent variables are binaries (coded 1 if the individual voted for the party and 0 if not) the coefficients can be interpreted as the marginal change in the probability of voting for the party. Other things being equal, therefore, the probability of voting Labour was a sizeable 0.11 (or about 11 percentage points) higher for an individual who reported being contacted by the party than for someone who did not. Being contacted by the Liberal Democrats had a similarly large effect on the probability of voting for that party, increasing by 0.09 (c. 9 percentage points) compared to those not contacted. The effect of being contacted by the Conservatives, meanwhile, was smaller, only increasing the probability of voting for that party by 0.05 (5 percentage points).

This would seem to imply that Labour's local campaign was the most effective in 2010. In a sense, it was, as the party got a better return to its efforts than its rivals (see Fisher et al, 2011). Even so, Labour lost the 2010 election badly. Its local campaign effort clearly could not counteract the other very substantial handicaps Labour faced in 2010 (a seriously weakened economy, a deeply unpopular leader, and so on): at best, it protected the party from an even worse defeat. Furthermore, compared to previous elections when New Labour was in the ascendant, the party's campaign resources were limited, and it was outspent by the Conservatives for the first time since the early 1990s (Johnston *et al*., 2011). Its campaign may have been more effective at the margin than the other parties' efforts, therefore, but with restricted resources available, there were limits to how far this could help the party. (It is also possible that, given its straightened circumstances, the party had to be more efficient with what it had available: had it been better resourced, it might not have been as focussed or as effective in its campaign efforts.)

Tackling selection bias through using DiD with interaction

While common and easy to implement, however, the conventional regression approach utilised above comes with a number of problems. Not least among them is the risk of selection bias. If self-reported contact with a party's campaign and vote choice are both influenced by a common factor – for instance, by pre-existing support for the party – there is a risk that the estimator for impact of contact on vote might be mis-specified. We need some

means of minimising this risk. The second difference-in-difference approach adopted here is in effect an extension of the regression approach to try and deal with this.[5]

At its core, DiD with interaction compares change over time in the behaviour being examined for those who are, and those who are not, exposed to some sort of intervention. Luckily, the 2010 BES Campaign Internet Panel survey is well-suited to DiD with interaction estimation given that we know vote intention at the start of the campaign and actual vote at the election four weeks later on for individuals who are contacted by each party and for those who are not. An adaptation of the regression method can then be used to capture a purer DiD with interaction estimate of the campaign effect. To model the effect of the change each individual is treated as two separate cases: a pre-treatment case and a post-treatment case. In addition to dummy variables for the outcome of interest and for exposure to the treatment a time variable is also included. By specifying an interaction effect between time and treatment it is possible to estimate the effect of exposure to the treatment on voting net of other influences. In particular, the presence of the treatment dummy acts to partial out potential biases around the types of voters that were contacted or around contact recall, leaving a purer estimator of the impact of that contact in the DiD with interaction estimator (the time-intervention interaction term).

To put this into effect here each BES panel respondent provides two cases. The first case records that individual's position in the pre-election wave, about a month before election day. In that wave individuals were asked whether they had decided how they would vote in the upcoming election. Respondents were asked whether they had definitely decided to vote for a party, or were leaning to it, though not yet firmly committed. For our purposes only definite intentions are treated as pre-election support for the party at the start of the campaign period and 'leaners' are grouped with those who were then thinking of voting for another party or intended to abstain. We use this expressed intention as the pre-election score for the dependent variable, vote. The second case for each respondent records his or her reported vote, with the dependent variable coded to indicate actual vote choice. In addition, each case has: a time variable (coded 0 for the cases at the start of the campaign and 1 for the post-election cases); dummy variables for exposure to each of the three main parties' face-to-face campaigns (coded 1 if the respondent reported being contacted by the party during the campaign and 0 otherwise: hence in all cases these variables are coded 0 for cases at the start of the election campaign); interactions between the time and campaign contact variables; and a series of explanatory variables. As noted above, these models are implemented using OLS and we are particularly interested in the coefficients for the interaction terms since these are the DiD with interaction estimates of the impact of contact on voting.

The results of those DiD with interaction models are shown in table 4. On the whole, the explanatory variables show basically the same relationships with vote as was the case in the standard models reported in Table 3. Turning to the variables for face-to-face contact the first thing to notice is that, as in the conventional regression analysis, the direct effects of campaign contact on vote are significant and positive in all three equations: those who recalled being contacted by a party were more likely to report either intending to vote for it before the election or actually doing so on polling day than were those who did not recall being contacted.

[5] See Ashenfelter and Card (1985) and Card and Kreuger (1994)

However, remember that each equation also contains a dummy variable for time (coding whether the data are from the start of the campaign or immediately after the election) and an interaction between the time and campaign contact variables. The direct effect of campaign contact therefore applies only to the cases from the start of the campaign period (where the variable for time and hence also the interaction term take values of 0). What the coefficients for campaign contact show, therefore, is that reporting being contacted by a party during the campaign is associated with an intention to vote for it at the start of the campaign. Clearly, since reported contact is based on post-election responses while vote intention at the start of the campaign was measured before the start of the campaign, the causal connection here is not straightforwardly from the 'explanatory' to the dependent variable: a simple "time's arrow" logic means that the future cannot cause the past, so vote intention when the campaign began cannot be caused by later remembering being contacted during the campaign. But this does suggest that, either because parties disproportionately target those voters already leaning towards them in order to mobilise their supporters, or because respondents are more likely to remember being contacted by parties they already support than by parties they do not, part of the apparent influence of party contact on vote can be accounted for by the association between intentions before the campaign and recalled interaction with the campaign. This may, of course, reflect longer-term local campaigning by parties in the months preceding the calling of the election, and there is evidence suggesting that such campaigning does build up support for parties (Cutts, 2006; Johnston et al., 2011; Cutts et al., 2012). However, this is not in itself evidence that face-to-face contact during the campaign makes a difference.

To get at that, we need to look at the interactions between face-to-face contact and time. These reveal a strikingly different picture to that implied by the conventional regression approach. Two of the interaction coefficients, for face-to-face contact by the Labour campaign in the Labour vote model and for contact by the Liberal Democrats in the Liberal Democrat vote model, are significant at the 0.01 level and in the expected direction. Labour and Liberal Democrat face-to-face efforts did win them extra support during the campaign, over and above what they might have expected before the election. But the equivalent interaction term in the Conservative model does not quite reach conventional levels of statistical significance ($p = 0.069$), though it is correctly signed. In other words, once we partial out the general drift of each party's support over the course of the campaign using this DiD with interaction method only the Labour and Liberal Democrat face-to-face campaigns emerge clearly as electoral assets for their parties.

The effect sizes are noteworthy too, being generally somewhat more modest than was the case in the conventional regression approach, though still respectable. Being contacted by Labour or the Liberal Democrats raised an individual's probability of voting for each party by around 0.08 (8 percentage points) compared to a similar individual who was not contacted.

Running a quasi-experiment: propensity score matching

Finally, what can we learn from the third strategy examined here, propensity score matching (PSM: Rosenbaum and Rubin, 1983; Dehejia and Wahba, 2002; Bryson et al., 2002; Jalan and Ravaillon, 2003)? This is a rather different approach to those described above. The basic idea is to replicate, *post hoc*, a random experimental design by matching each individual in a survey who did receive the treatment of interest with an observably similar individual who did not receive that treatment but who could well have done so. In our example, exposure to each party's face-to-face campaign will never be universal and parties concentrate their efforts on contacting individuals who are liable to vote for them or have a reasonable

likelihood of being persuaded to do so, especially if they live in marginal seats which could easily change hands. They will tend not to expend energy contacting other voters who are unlikely to ever support them. In consequence, party campaigns will be targeted geographically (on the marginal seats) and at the level of individuals (as far as possible, at actual or potential supporters). But even the best-organised and resourced campaign will be unable to contact every individual they may wish to reach. Hence there will be individuals who have many of the characteristics of someone the party might also wish to contact but who will not, in fact, be contacted. PSM identifies such individuals and matches them with ostensibly similar individuals who were contacted (or, in our case, who reported being contacted). Given that matched treatment and control individuals are, by definition, designed to be similar (though how similar depends on the effectiveness of the matching) the effect of the intervention is assumed to be the difference in voting behaviour between the matched treatment and control groups.

To achieve this matching, PSM models exposure to the treatment (here, contact by a campaign).[6] For each individual, we use logit models to predict their probability of exposure to the campaign (their propensity score). Individuals who did report contact are matched with individuals who did not report contact but whose propensity scores for contact are as similar as possible to their own using single nearest neighbour matching (i.e. matching each contacted individual to the single uncontacted individual with the nearest propensity score). This matched sample is similar to a randomised experimental design, in so far as the underlying characteristics of the treatment and control groups are similar, with the main exception being whether or not they were exposed to a party's face-to-face campaign. That said, the matching process is probabilistic, not exact[7] and based inevitably only on observable characteristics available in the data (and hence vulnerable to omitted variable biases). The harder it is to model the likelihood of receiving the intervention therefore – which in our case is whether the individual was contacted or not – then the larger the margin of error in the matching process is likely to be (for a critique of propensity score methods, see e.g. Arceneaux *et al*., 2006, 2010).

Table 5 reports the logit models predicting who reported being contacted by each party's campaign. As expected, those already inclined towards each party (whether as party identifiers or past voters) were more likely to recall being contacted by it than were those who were not so inclined. What is more, wider engagement with the political process mattered. Respondents who had not sought help from their MP (whichever party that MP represented) were less likely to report being contacted by each party's campaign during the 2010 election than were respondents who had contacted their MP (whether because the latter were more interested in politics than the former, or because they were more likely to be on parties' databases of voters to contact). And the more attention individuals claimed they paid to politics at the start of the campaign the more likely they were to report being contacted by each party during the campaign. Finally, respondents were more likely to report being contacted in constituencies where the parties campaigned hard than where they made little effort: the more each party spent on its constituency campaign in each seat, the more likely individuals living there were to report being contacted by that party. And there is evidence that the harder a party's rivals worked in a seat, the more likely voters living there were to report being contacted by that party. The more the Liberal Democrats spent on their local campaign, for instance, the more likely respondents were to report being contacted not just by

[6] To conduct the PSM we make use of the psmatch2 command in Stata (Leuven and Sianesi, 2003).
[7] Nor, indeed, random as in a true randomised control trial.

the Liberal Democrats but also by Labour and the Conservatives, who were countering their opponent's campaign. Similarly, the harder Labour campaigned, the more likely respondents were to recall being contacted by the Conservatives.

These models are used to estimate a propensity score for each respondent, which is then used to match contacted individuals with others who had a similar chance of being contacted given their characteristics but who did not report being contacted. By comparing voting behaviour across the matched treatment and control groups we can obtain an alternative measure for the impact of face-to-face canvassing on party support. These results suggest a remarkably consistent effect (table 6). Among these matched pairs, being contacted by a party raises the chances of voting for that party by between 6 and 7 percentage points over the chances of voting for it having not been contacted. Clearly, this is a much smaller effect than the crude estimate provided in Table 1. But it is far from negligible and in a close contest could make all the difference between winning and losing a seat.

But PSM's effectiveness depends on effectively modelling campaign contact whether in order to create substantively meaningful matches. However, the fit statistics of our models predicting contact are relatively poor: much remains unexplained in terms of who is contacted.

Comparing effect sizes

How do the different methodologies and their respective results compare? Figure 2 shows the estimated effect sizes from each method, with their associated 95% confidence intervals. To give a bench-mark, the first estimate in each graph shows the raw difference in voting for each party between those who do and those who do not report being contacted by the party (i.e. before any statistical controls are applied). Not surprisingly, these are considerably (and significantly: see table 7) larger, by factors of three or more, than the effect sizes obtained from the various modelling exercises. As expected, the raw estimates very seriously exaggerate the likely impact on vote choice of being contacted by a campaign.

We are more interested, however, in how the different estimation methods compare. There is a relatively clear story here too, in that the estimates from the DiD with interaction models are consistently smaller than those for either conventional regression or PSM. The latter two tend to be quite similar in size for all three parties. But the real story is contained in the confidence intervals for the modelled estimates, all of which overlap with each other to a substantial extent. In other words (and as table 7 confirms more rigorously), the modelled estimates are, on the whole, statistically indistinguishable from each other: such differences as do appear are generally within the margin of error. The only difference which reaches 95% significance (and then only just) is that between the conventional regression and the DiD with interaction estimates for the Conservatives: the former effect is slightly higher than the latter. It is worth noting, however, that the equivalent comparison between the conventional regression and DiD with interaction estimates for Labour falls just short of significance at the 95% level. Despite anxieties over their reliability, therefore, the estimates derived from the conventional regression approach stand comparison with more carefully constructed estimates using DiD with interaction methods.

Conclusions

Where does this leave us? Certainly, all the methods examined here confirm once again the importance of local campaign effort in modern elections. Where parties work hard, they tend to gain electoral dividends. In particular, individuals who are contacted by parties are more likely to vote for them than are individuals who are not. The analyses reported here confirm that this is unlikely to be an artefact of survey response biases: no matter how we look at the data, the same results recur, giving us confidence in their validity. However, the parties did not all receive the same rewards for their efforts. The Liberal Democrats and especially Labour obtained greater returns from their campaign contacts with voters than did the Conservatives. In part, this may reflect the long-standing observation that in recent elections the Conservatives tend to get fewer returns to their local campaigning than the other parties (e.g. Pattie *et al*., 1995; Denver *et al*., 2002). But as discussed above, it also reflects the parties' very different positions in the run-up to the 2010 elections (see also Fisher *et al*., 2011). The Conservatives were by far the best-resourced party of the three, but had a substantial job to do to win sufficient seats to form a government (a feat which would have required one of the largest swings from government to opposition of modern times). Their resources were spread widely, therefore. Labour and the Liberal Democrats, working in more straightened circumstances, had to focus their campaign resources in a relatively few places: that concentration seems to have helped them do relatively well where they were able to put up a strong local fight. A relatively effective constituency campaign could not compensate Labour for its failings in office. But it does seem to have helped Labour contain its losses to manageable proportions. The party would almost certainly have gone down to an even more serious loss of its parliamentary base had its local campaigns been less effective.

That the three estimation methods employed here – the conventional regression, DiD with interaction and PSM – all broadly agree is reassuring, implying as it does that estimates from the conventional regression approach so widely employed in the literature are not too far from the mark. Even so, the results discussed above should not encourage complacency. There is no guarantee that the various methods used here would necessarily always agree so well in other elections. Furthermore, the well-known methodological problems inherent in the conventional regression approach and outlined at some length earlier in the paper cannot just be wished away. In the absence of a genuine experiment, DiD with interaction is more likely to produce a reliable estimate of campaign effects than conventional regression. Like Arceneaux *et al*. (2006, 2010), however, we are more sceptical about the utility of PSM methods here in the absence of strong models of reported party contact. We would therefore recommend DiD with interaction. These estimates are more conservative than the others, making this approach perhaps the most desirable for the cautious analyst. More importantly, it is more likely than is conventional regression to give a relatively robust estimate of campaign effect sizes.

References

Ai, C. and Norton, E. 2003. Interaction terms in logit and probit models. *Economics Letters*, 80, 123-129

Ansolabehere, S. and Iyengar, S. 1995. *Going Negative: How Political Advertisements Shrink and Polarize the Electorate*. New York: The Free Press.

Arceneaux, K, Gerber, A.S. and Green, D.P. 2006. Comparing experimental and matching methods using a large-scale voter mobilization experiment. *Political Analysis*, 14, 37-62.

Arceneaux, K., Gerber, A.S. and Green, D.P. 2010. A cautionary note on the use of matching to estimate causal effects: an empirical example comparing matching estimates to an experimental benchmark. *Sociological Methods and Research*, 39, 256-282.

Ashenfelter, O. And Card, D. 1985. Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics*. 67, 648-660.

Bochel, J. and Denver, D. 1971. Canvassing, turnout and party support: an experiment. *British Journal of Political Science*. 1, 257-269.

Bryson, A., Dorsett, R. and Purdon, S. 2002. *The use of propensity score matching in the evaluation of active labour market policies*. London: Department for Work and Pensions.

Card, D. and Kreuger, A. 1994. Minimum wages and employment: a case study of the fast food industry in New Jersey and Pennsylvania. *American Economic Review*. 84(4), 772-784

Clarke, H.D., Sanders, D., Stewart, M.C. and Whiteley, P.F., 2004. *Political Choice in Britain*. Oxford: Oxford University Press.

Clarke, H.D., Sanders, D., Stewart, M.C. and Whiteley, P.F., 2009. *Performance Politics and the British Voter*. Cambridge: Cambridge University Press.

Cutts, D. 2006. Continuous activism and electoral outcomes: the Liberal Democrats in Bath. *Political Geography*. 25: 72-88.

Cutts, D., Johnston, R. J., Pattie, C. J., and Fisher, J., 2012. Laying the foundations for electoral success: Conservative pre-campaign canvassing before the 2010 British general election. *Journal of Elections, Public Opinion and Parties*, 22, 359-375.

Davenport, T.C., Gerber, A.S., Green, D.P., Larimer, C.W., Mann, C.B. and Panagopoulos, C. 2010. The enduring effects of social pressure: tracking campaign experiments over a series of elections. *Political Behavior*. 32, 423-430.

Dehejia, R. and Wahba, S. 2002. Propensity score matching methods for nonexperimental causal studies. *The Review of Economics and Statistics*. 84(1), 151-161

Denver, D., Hands, G., Fisher, J. and MacAllister, I. 2002. The impact of constituency campaigning in the 2001 General Election. In Bennie, L., Rallings, C., Tonge, J. and Webb, P. (eds) *British Elections and Parties Review 12: The 2001 General Election*. London: Frank Cass.

Denver, D., Hands, G. and MacAllister, I. 2004. The electoral impact of constituency campaigning in Britain, 1992-2001. *Political Studies.* 52, 289-306.

Farrell, D.M. and Webb, P. 2002. Political parties as campaign organisations. In Dalton, R.J. and Wattenburg, M.P. (eds) *Parties Without Partisans*. Oxford: Oxford University Press.

Fisher, J. Cutts, D. and Fieldhouse, E. 2011. The electoral effectiveness of constituency campaigning in the 2010 British general election: the 'triumph' of Labour? *Electoral Studies*, 30, 816-828.

Fisher, J. and Denver, D. 2008. From foot-slogging to call centres and direct mail: a framework for analysing the development of district-level campaigning. *European Journal of Political Research*. 47: 794-826.

Fu, A., Dow, W. And Liu, G. 2007. Propensity score and difference-in-difference methods: a study of second-generation antidepressant use in patient with bipolar disorder. *Health Services and Outcomes Research Methodology, 7(1-2), 23-38*

Gerber, A.S., Green, D.P. and Larimer, C.W. 2008. Social pressure and voter turnout: evidence from a large-scale field experiment. *American Political Science Review*. 102, 33-48.

Gerber, A.S., Green, D.P. and Larimer, C.W. 2010. An experiment testing the relative effectiveness of encouraging voter participation by inducing feelings of pride or shame. *Political Behavior*. 32, 409-422.

Green, D.P. and Gerber, A.S., 2004. *Get Out the Vote: How to Increase Voter Turnout*. Washington D.C.: Brookings Institute.

Ikenwilo, D. 2013. A difference-in-difference analysis of the effect of free dental check-ups in Scotland. *Social Science and Medicine*, 83, 10-18.

Jalan, J. and Ravallion, M. 2003. Estimating the benefit incidence of an antipoverty program by propensity score matching. *Journal of Business and Economic Statistics*. 21(1): 19-30

Johnston, R. J., Cutts, D., Pattie, C. J. and Fisher, J. 2012. We've got them on the list: canvassing and voting in a British general election campaign. *Election Studies*, 31, 317-329.

Johnston, R.J., Pattie, C.J., Cutts, D., Fieldhouse, E., and Fisher, J. 2011. Local campaign spending at the 2010 General Election and its impact: exploring what wider regulation has revealed. *Political Quarterly*. 82, 169-192.

Johnston, R.J., Pattie, C.J., Cutts, D. and Fisher, J., 2012. Spending, contacting and voting: the 2010 British general election in the constituencies. *Environment and Planning A*. 44, 1165-1184.

Johnston, R.J., Pattie, C.J., Fisher, J., Cutts, D. and Fieldhouse, E. 2014. The long and the short of it: local campaigning at the British 2010 General Election. *Political Studies*, 61(S1), 114-137.

John, P. and Brannan, T. 2008. How different are telephoning and canvassing? Results from a 'get out the vote' field experiment in the British 2005 General Election. *British Journal of Political Science*. 34: 334-362.

Liu, X., Yan, H. And Wang, D. 2010. The evaluation of "Safe Motherhood" program on maternal care utilization in rural western China: a difference in difference approach. *Biomedical Central Public Health*. 10: 566.

Leuven, E. and Sianesi, B. 2003. PSMATCH2: Stata module to perform full Mahalnobis and propensity score matching, common support programming, and covariate imbalance testing. http://ideas/repec.org/c/boc/bocode/s432001.html. This version 4.0.6.

Machin, S., McNally, S. and Meghir, C. 2004. Improving pupil performance in English secondary schools: Excellence in Cities. *Journal of the European Economic Association*. 2(2-3), 396-405

Meghir, C. and Palme, M. 2005. Educational reform, ability and family background. *The American Economic Review*. 95(1), 414-424

Norris, P., Curtice, J., Sanders, D., Scammell, M. and Semetko, H.A. 1999. *On Message: Communicating the Campaign*. London: Sage.

Panagopoulos, C. 2010. Affect, social pressure and prosocial motivation: field experimental evidence of the mobilizing effects of pride, shame and publicizing voting behavior. *Political Behavior*. 32, 369-386.

Pattie, C.J. and Johnston, R.J. 2003. Local battles in a national landslide: constituency campaigning at the 2001 British General Election. *Political Geography*. 36, 381-414.

Pattie, C.J. and Johnston, R.J. 2010. Constituency campaigning and local contests at the 2010 UK General Election. *British Politics*. 5, 481-505.

Pattie, C.J., Johnston, R.J. and Fieldhouse, E.A. 1995. Winning the local vote: the effectiveness of constituency campaign spending in Great Britain. *American Political Science Review*. 89, 969-983.

Pattie, C.J., Whiteley, P., Johnston, R. and Seyd, P., 1994. Measuring local campaign effects: Labour party constituency campaigning at the 1987 General Election. *Political Studies*. 42, 469-479.

Rosenbaum, P.R. and Rubin, D.B., 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41-55.

Vavreck, L. 2007. The exaggerated effects of advertising on turnout: the dangers of self-reports. *Quarterly Journal of Political Science*, 2, 325-343.

Weinick, R., Zuvekas, H. And Cohen, J. 2000. Racial and ethnic differences in access to and use of health care services, 1977 to 1996. *Medical Care Research and Review*, 57, 36-54

Wilkinson, K. and McLennan, D. 2010. *Narrowing the gap? Analysing the impact of the New Deal for Communities on educational attainment.* London: Department for Communities and Local Government

Table 1: The impact of face-to-face contact in the 2010 election campaign: a first cut with no statistical controls (source: 2010 BES Campaign Internet Panel)

Party	Conservatives	Labour	Liberal Democrat
Contacted by party: % voting for party	52.7	53.1	49.8
Not contacted by party: % voting for party	30.8	21.4	23.3
Difference	21.9	31.6	26.5
SE	1.3	1.4	1.5
T	16.65**	23.12**	17.25**
N	12598	12598	12598

** Significant at p = 0.01

Table 2: Pre-election party identification and self-reported experience of face-to-face contact by party campaigns, 2010 (source: British Election Study 2010 Campaign Internet Panel)

	Contacted face-to-face by party?					
	Conservative		Labour		Liberal Democrat	
Face-to-face contact reported?	No	Yes	No	Yes	No	Yes
	%	%	%	%	%	%
Pre-election party identification:						
Conservative	27.8	46.1	31.1	14.4	30.2	22.0
Labour	33.9	24.8	30.7	56.9	33.3	27.7
Liberal Democrat	15.8	14.6	16.0	12.2	14.4	34.4
Other	11.5	7.7	11.2	9.6	11.2	8.5
None/DK	11.1	6.8	11.0	6.9	10.9	7.3
N	11617	1367	11904	1083	12139	845

Table 3: Predicting the impact of self-reported face to face contact on voting at the 2010 election: conventional regression OLS models with robust standard errors (source: 2010 BES Campaign Internet Panel)

	Vote 2010		
	Labour	Conservative	Liberal Democrat
Vote, 2005 (comparison = did not vote)			
Labour	0.1793**	0.0505**	0.0422**
Conservative	-0.0052	0.2872**	-0.0143
Lib Dem	-0.0455**	0.0235+	0.3038**
Other	0.0042	0.0366*	-0.0074
Too young	0.0613*	0.0472*	0.0628*
Don't know	0.0198	0.0937**	0.0806**
Party identification, pre-election (comparison = no party ID)			
Very strong Labour identification	0.5669**	-0.2011**	-0.1935**
Fairly strong Labour identification	0.4078**	-0.1571**	-0.0903**
Not very strong Labour identification	0.2095**	-0.0952**	-0.0017
Very strong Conservative identification	-0.0944**	0.4885**	-0.2089**
Fairly strong Conservative identification	-0.0875**	0.4702**	-0.1980**
Not very strong Conservative identification	-0.0919**	0.3591**	-0.1193**
Very strong Liberal Democrat identification	-0.0500**	-0.1805**	0.3960**
Fairly strong Liberal Democrat identification	-0.0431**	-0.1509**	0.3740**
Not very strong Liberal Democrat identification	-0.0242	-0.0973**	0.2469**
Other	-0.0399**	-0.0460**	-0.0854**
Respondent sought help from local MP (comparison = yes)			
No	-0.0042	0.0021	0.0053
Pre-election attention to politics (10=high)	-0.0006	0.0070**	0.0003
Labour short campaign spend %	0.0007**	-0.0000	-0.0005**
Conservative short campaign spend %	-0.0003**	0.0005**	0.0002
Lib Dem short campaign spend %	-0.0005**	-0.0002*	0.0008**
Labour face-to-face campaign contact	0.1141**		
Conservative face-to-face campaign contact		0.0526**	
Lib Dem face-to-face campaign contact			0.0935**
Constant	0.1102	0.0757	0.2098
R^2	0.4802	0.5623	0.3330
N	12155	12155	12155

\+ Significant at p=0.10
* Significant at p= 0.05
** Significant at p = 0.01

Table 4: Predicting the impact of self-reported face to face contact on voting at the 2010 election: difference-in-difference with time-intervention interaction OLS models with robust standard errors (source: 2010 BES Campaign Internet Panel)

	Vote 2010		
	Labour	Conservative	Liberal Democrat
Vote, 2005 (comparison = did not vote)			
Labour	0.1457**	0.0395**	0.0138
Conservative	-0.0001	0.2336**	-0.0104
Lib Dem	-0.0358**	0.0157+	0.2388**
Other	-0.0035	0.0179+	-0.0116
Too young	0.0605**	0.0302*	0.0508**
Don't know	0.0047	0.0323**	0.0325**
Party identification, pre-election (comparison = no party ID)			
Very strong Labour identification	0.6716**	-0.1485**	-0.0882**
Fairly strong Labour identification	0.4833**	-0.1138**	-0.0273*
Not very strong Labour identification	0.2278**	-0.0607**	0.0091
Very strong Conservative identification	-0.0568**	0.5953**	-0.1106**
Fairly strong Conservative identification	-0.0517**	0.5537**	-0.1059**
Not very strong Conservative identification	-0.0553**	0.3744**	-0.0666**
Very strong Liberal Democrat identification	-0.0188+	-0.1307**	0.5670**
Fairly strong Liberal Democrat identification	-0.0166+	-0.1067**	0.4781**
Not very strong Liberal Democrat identification	-0.0042	-0.0715**	0.2460**
Other	-0.0199*	-0.0382**	-0.0442
Respondent sought help from local MP (comparison = yes)			
No	0.0028	0.0017	-0.0010
Pre-election attention to politics (10=high)	0.0001	0.0052**	-0.0002
Labour short campaign spend %	0.0004**	-0.0001	-0.0003**
Conservative short campaign spend %	-0.0002*	0.0003**	0.0001
Lib Dem short campaign spend %	-0.0002**	-0.0001*	0.0005**
Labour face-to-face campaign contact	0.0386**		
Conservative face-to-face campaign contact		0.0286**	
Lib Dem face-to-face campaign contact			0.0281*
Time post-election (comparison=pre-election)	0.0308**	0.0503**	0.1412**
Time*Lab face-to-face campaign contact	0.0772**		
Time*Con face-to-face campaign contact		0.0233+	
Time*LD face-to-face campaign contact			0.0794**
Constant	0.0359	0.0389	0.0483
R^2	0.5108	0.5869	0.3766
N	23410	23410	23410

\+ Significant at p=0.10
* Significant at p= 0.05
** Significant at p = 0.01

Table 5: Propensity score matching stage 1: predicting who should have been contacted face to face by parties in the 2010 campaign (logit models. Source: 2010 BES Internet Campaign Panel)

	Contacted face-to-face by:		
	Labour	Conservative	Liberal Democrat
Vote, 2005 (comparison = did not vote)			
Labour	0.390	0.088	0.229
Conservative	0.026	0.401**	0.074
Lib Dem	-0.169	0.164	0.848**
Other	0.108	0.064	0.217
Too young	0.017	0.280	0.481
Don't know	0.044	0.033	0.578*
Party identification, pre-election (comparison = no party ID)			
Very strong Labour identification	1.234**	-0.083	0.212
Fairly strong Labour identification	0.565**	0.145	0.393*
Not very strong Labour identification	0.565**	0.145	0.215
Very strong Conservative identification	-0.470	0.894**	0.206
Fairly strong Conservative identification	-0.332	0.655*	0.066
Not very strong Conservative identification	-0.152	0.436**	0.164
Very strong Liberal Democrat identification	-0.109	-0.353	1.707**
Fairly strong Liberal Democrat identification	0.080	0.054	0.746**
Not very strong Liberal Democrat identification	0.244	0.282	0.642**
Other	0.165	-0.020	0.014
Respondent sought help from local MP (comparison = yes)			
No	-0.362**	-0.127*	-0.148*
Pre-election attention to politics (10=high)	0.058**	0.098**	0.062**
Labour short campaign spend %	0.019**	0.004**	0.002
Conservative short campaign spend %	-0.000	0.013**	0.002
Lib Dem short campaign spend %	0.004**	0.003**	0.017**
Constant	-4.187	-4.416	-4.669
-2 log likelihood			
Improvement	754.39	517.94	629.44
Sig	0.000	0.000	0.000
% correctly classified			
Nagelkerke R^2	0.107	0.059	0.102
N	12598	12598	12598

\+ Significant at p=0.10
* Significant at p= 0.05
** Significant at p = 0.01

Table 6: Propensity score matching results: the effect of face-to-face contact on voting for a party in 2010

Party	Conservatives	Labour	Liberal Democrat
Contacted by party: % voting for party	52.7	53.1	49.8
Not contacted by party: % voting for party	47.7	42.7	40.0
Difference	5.0	10.4	9.8
SE	2.0	2.4	2.6
T	2.44*	4.38**	3.80**
N	2820	2028	1670

* Significant at p= 0.05
** Significant at p = 0.01

Table 7: Comparing estimates: t-tests

t-values

a) Labour estimates:

	No control	Conventional OLS	DiD with interaction
Conventional OLS	10.972**		
DiD with interaction	11.393**	1.896+	
PSM	7.630**	0.388	-0.936

b) Conservative estimates:

	No control	Conventional OLS	DiD with interaction
Conventional OLS	10.259**		
DiD with interaction	10.727**	1.983*	
PSM	7.085**	0.117	-1.124

c) Liberal Democrat estimates

	No control	Conventional OLS	DiD with interaction
Conventional OLS	8.031**		
DiD with interaction	7.717**	0.583	
PSM	5.564**	-0.149	-0.715

\+ Significant at p=0.10
* Significant at p= 0.05
** Significant at p = 0.01

Figure 1: A hypothetical illustration of difference-in-difference estimation

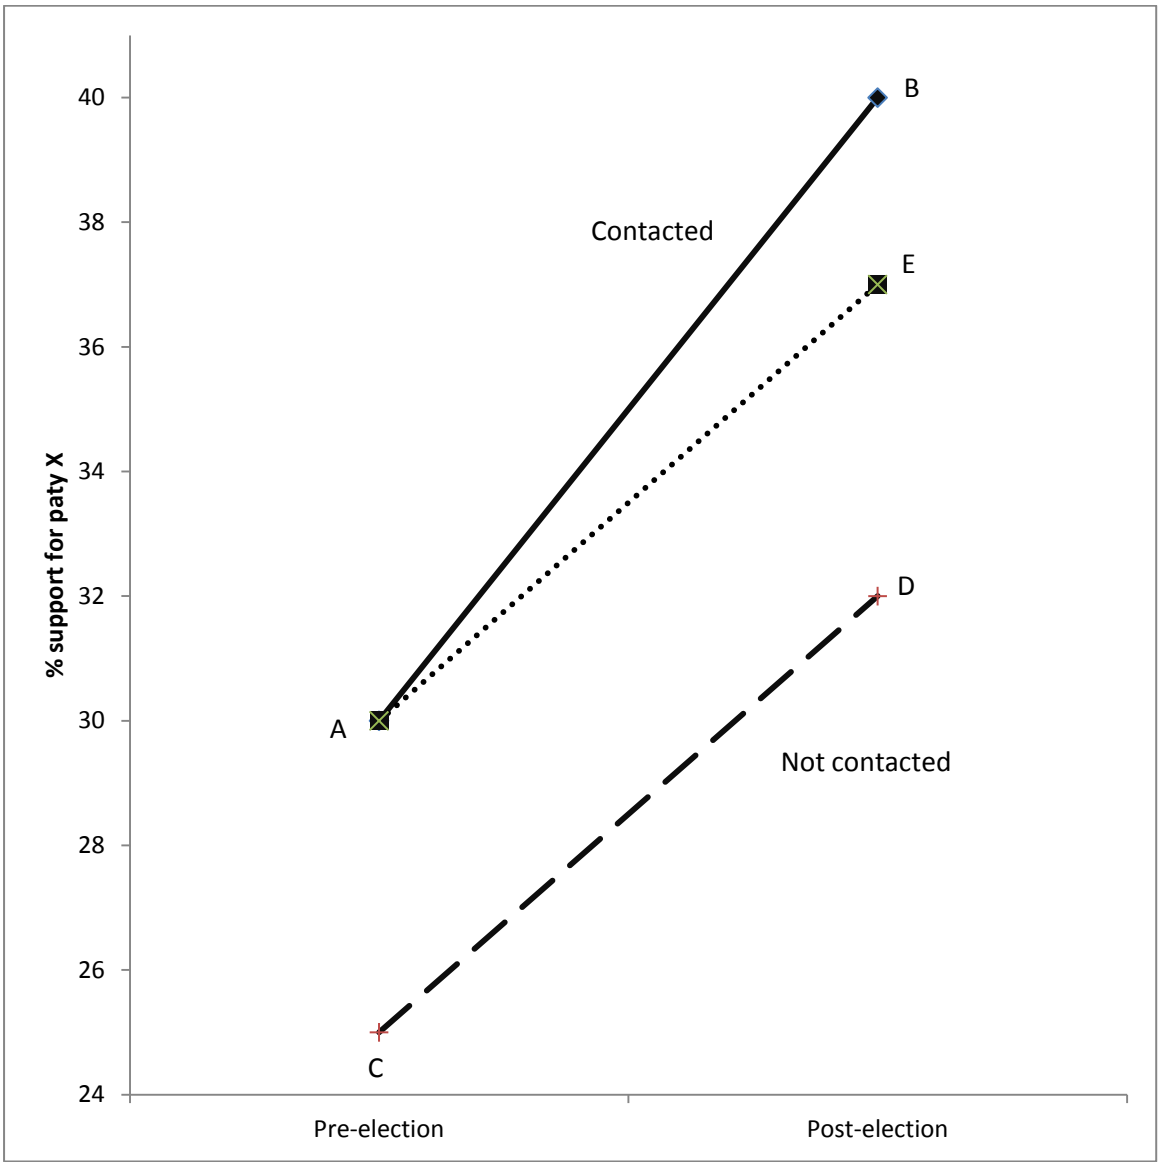

Figure 2: Comparing effect sizes across different estimation approaches

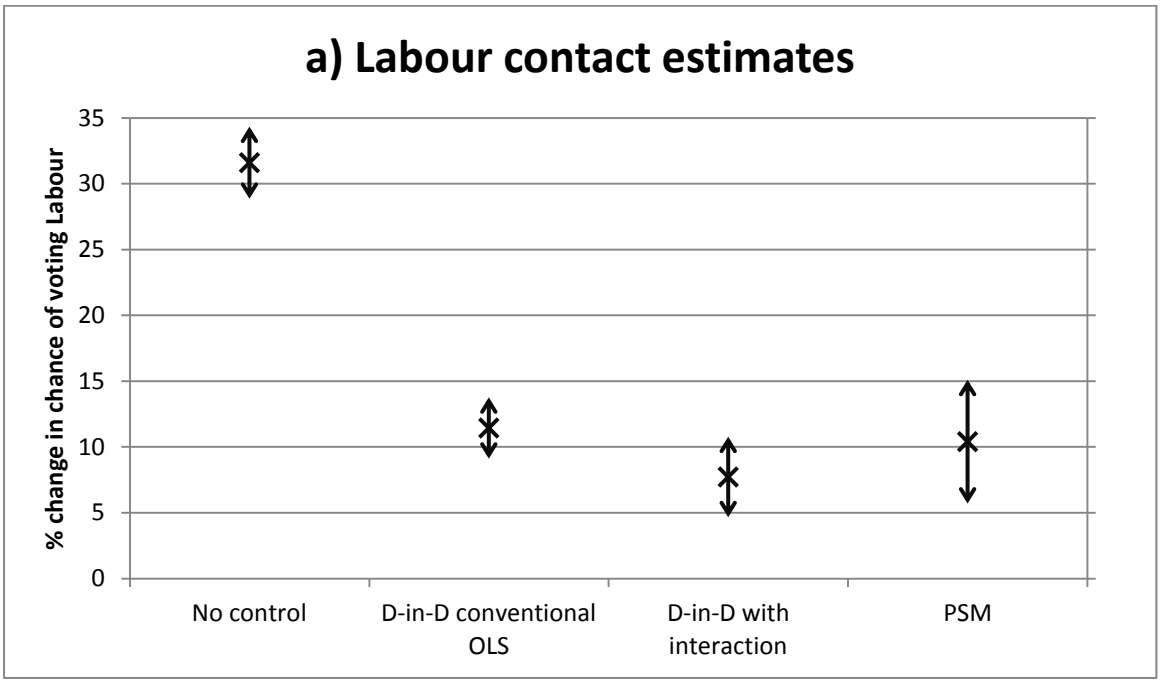

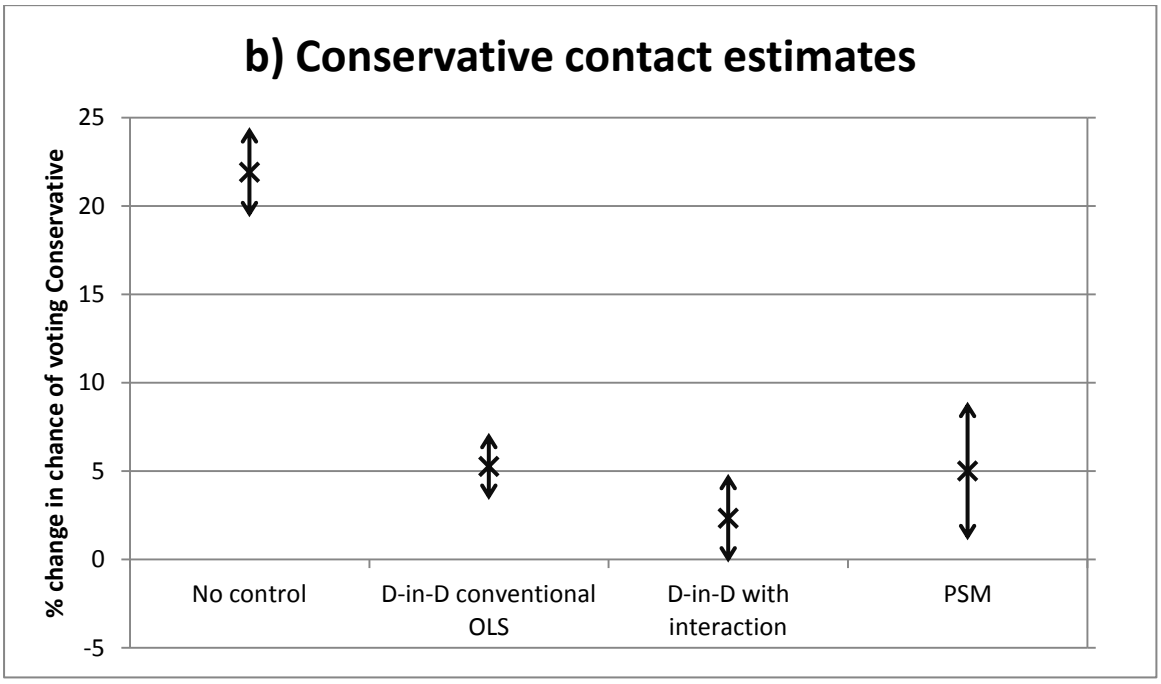

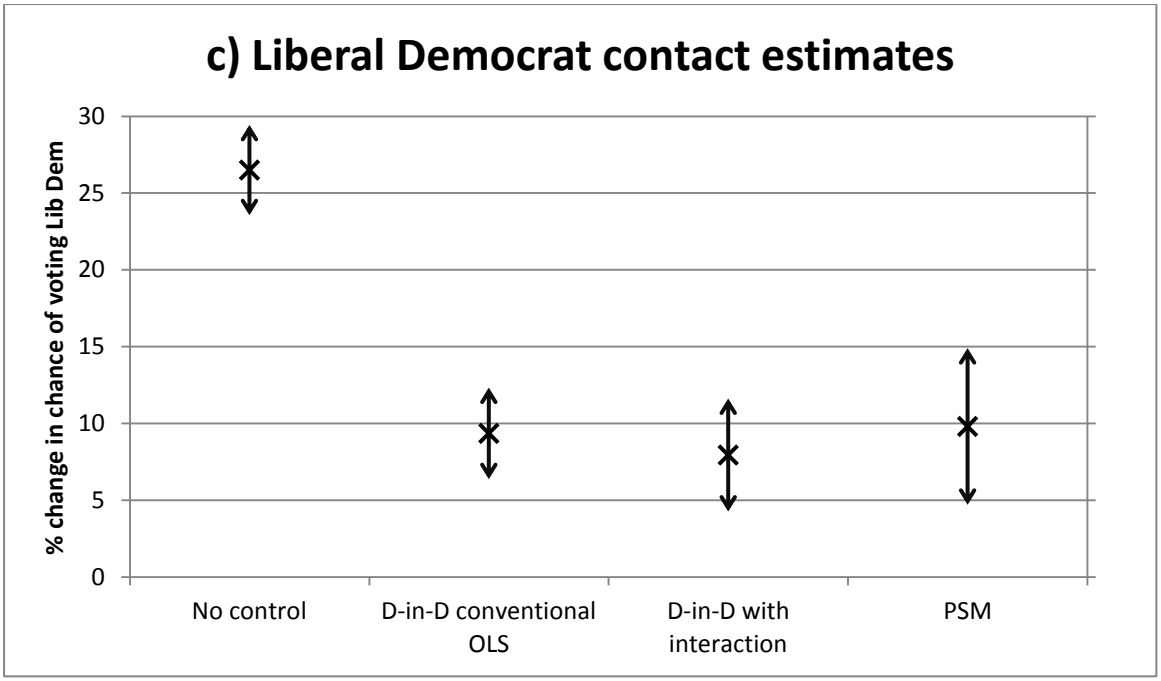
c) Liberal Democrat contact estimates
% change in chance of voting Lib Dem
30
25
20
15
10
5
0
No control
D-in-D conventional OLS
D-in-D with interaction
PSM