Online Appendix for *Who Votes More Strategically?*[1]

Andrew C. Eggers[†]

Nick Vivyan[‡]

[1] This version: January 24, 2020.

[†] Nuffield College and Department of Politics and International Relations, University of Oxford, UK. email: andrew.eggers@nuffield.ox.ac.uk

[‡] School of Government and International Affairs, Durham University, UK. email: nick.vivyan@durham.ac.uk

Appendix A: Details on estimating tactical incentives

The tactical incentive τ is defined as the maximum difference in expected utility between an insincere vote and a sincere vote:

$$\tau \equiv \max_j \mathbf{p}(j) \cdot \mathbf{u} - \mathbf{p}(1) \cdot \mathbf{u}, \tag{A.1}$$

where $\mathbf{p}(j) \equiv \{p_1(j), p_2(j), \ldots, p_K(j)\}$ is the probability of each candidate being elected when the voter votes for candidate j and $\mathbf{u} = \{u_1, u_2, \ldots, u_K\}$ is the voter's VNM utility from electing each of the K candidates. To measure τ it is useful to decompose $\mathbf{p}(j)$ as follows. Let π_k denote the probability that candidate k wins by more than 1 vote, and let π_{jk} denote the probability that candidate j and k tie for first. Assume that (i) for any pair of candidates there is a fixed rule about which one wins in the event of a tie,[2] (ii) the probability of any candidate leading another by one vote is the same as the probability of those two candidates being tied for first, and (iii) ties involving more than two candidates can be ignored (as Fisher and Myatt (2017) show is reasonable in a large electorate). Then $p_j(j) = \pi_j + 2\sum_{k \neq j} \pi_{jk}$ and $p_k(j) = \pi_k + \sum_{l \notin \{j,k\}} \pi_{kl}$, so that e.g.

$$\mathbf{p}(1) = \left\{ \pi_1 + 2\sum_{k \neq 1} \pi_{1k}, \pi_2 + \sum_{l \notin \{1,2\}} \pi_{2l}, \pi_3 + \sum_{l \notin \{1,3\}} \pi_{3l}, \ldots, \pi_K + \sum_{l \notin \{1,K\}} \pi_{Kl} \right\}.$$

Now, define $\tilde{\mathbf{p}}(j) = \mathbf{p}(j) - \{\pi_1, \pi_2, \ldots, \pi_K\}$ and observe that $\tilde{\mathbf{p}}(j) \cdot \mathbf{u} = \mathbf{p}(j) \cdot \mathbf{u} - C$, from which it follows that we can estimate τ using the $\tilde{\mathbf{p}}(j)$s instead of $\mathbf{p}(j)$s. This means that we can estimate τ as a function of pivot probabilities and utilities only (ignoring events in which any candidate wins by more than one vote), which is helpful because it is easier to estimate e.g. π_{jk} precisely than it is to estimate e.g. π_j precisely.

To estimate the pivot probabilities, we begin with a model of counterfactual election outcomes.[3] We model counterfactual candidate vote shares using a Dirichlet distribution (c.f. Fisher and Myatt, 2017), which assigns a positive probability mass to every point on a simplex. The distribution of vote shares for K parties can be characterized by a Dirichlet distribution with parameter vector $s\mathbf{v} \equiv \{sv_1, sv_2, \ldots, sv_K\}$, where $\mathbf{v}$ is the expected value of the distribution and s is a measure of precision. As noted in the main text of the article, we set $\mathbf{v}$ equal to the observed vote shares in the election. To ensure that our model has a level of uncertainty similar to that of an informed expert, we set s to maximize the joint likelihood of all constituency-level forecasts in the 2005, 2010, and 2015 elections assuming that the mean result in each constituency is the observed result. More specifically, we solve the problem

$$\arg\max_s \prod_t \prod_i \mathrm{Dir}\big(\mathbf{x}_{it}; s\mathbf{v}_{it}\big) \tag{A.2}$$

where s is the level of precision, $\mathbf{x}_{it}$ and $\mathbf{v}_{it}$ are the vector of forecasted vote shares and actual vote shares (respectively) in constituency i at time t, and $\mathrm{Dir}(\mathbf{x}_{it}; s\mathbf{v}_{it})$ gives the density of the Dirichlet distribution at $\mathbf{x}_{it}$ given parameter vector $s\mathbf{v}_{it}$.

The next step is to derive pivot probabilities from this model. Fisher and Myatt (2017) show that with a Dirichlet model of election outcomes for three candidates one can calculate pivotal probabilities analytically.[4] Here we develop a more flexible approach that can be applied to an

[2] This assumption is innocuous but eliminates the need to distinguish between ties and near-ties.

[3] In a similar exercise, Gelman, King and Boscardin (1998) and Gelman, Silver and Edlin (2012) use a model of U.S. presidential election outcomes to estimate the probability of casting a decisive vote in each state.

[4] See also Hoffman (1982), who discusses the general case and numerically computes pivot probabilities for

arbitrary number of candidates as well as alternative electoral rules.

A brute-force way to approximate the pivot probability for a pair of candidates is to draw a large number of simulated elections from the model and compute the proportion of ties for first involving that pair. (Given the Dirichlet is a continuous distribution, a "tie for first" could mean the two candidates are separated by a share of less than $1/n$, where n is the number of voters, and others get lower shares.) To improve precision for a given number of simulated elections, one might instead calculate the proportion of simulations in which the two candidates finish in the lead with vote shares within x of each other, and divide that proportion by xn; there is a bias-variance tradeoff involved, with larger x causing bias to the extent that the density differs at ties for first and near-ties for first.

We introduce a numerical alternative that generally reproduces the pivot probabilities estimated by the brute-force simulation approach but much more quickly, particularly when the pivot probabilities are very low. Given K parties, the pivot probability for parties 1 and 2 is approximately

$$\frac{1}{n}\int_{1/K}^{1/2} \Pr(x_1 = x_2 = y, x_3 < y, \dots, x_K < y)dy, \tag{A.3}$$

where $x_1, x_2, \dots x_K$ are realized vote shares for parties $1, 2, \dots K$. To see this, note first that we want to calculate the probability of a vector of vote shares where 1 and 2 are essentially tied and a single vote could change the outcome. The change in a candidate's vote share accomplished by adding a single vote is approximately $1/n$ (the leading term in expression A.3).[5] $\Pr(x_1 = x_2 = y, x_3 < y, \dots, x_K < y)$ describes the probability of candidates 1 and 2 tying for first with a vote share of y; we integrate for y between $1/K$ and $1/2$, which are the lowest and highest possible vote shares of two candidates tied for first. To make progress in computing expression A.3, we transform the joint probability into a conditional probability and make an independence assumption:

$$\begin{aligned}\Pr(x_1 = x_2 = y, x_3 < y, \dots, x_K < y) &= \Pr(x_1 = x_2 = y)\Pr(x_3 < y, \dots, x_K < y | x_1 = x_2 = y) \\ &\approx \Pr(x_1 = x_2 = y)\prod_{i=3}^{K}\Pr(x_i < y | x_1 = x_2 = y)\end{aligned} \tag{A.4}$$

The independence assumption we make is that the joint probability of all parties from 3 to K being below y (conditional on 1 and 2 get y) is given by the product of the conditional probabilities for each party taken separately. This assumption saves us from having to calculate an integral over multiple dimensions, and we show below that its implications appear to be fairly innocuous. Next we apply the "aggregate property" of Dirichlet distributions to compute this object for a specific model of election outcomes. If we have K vote shares distributed according to $\text{Dir}(\alpha_1, \alpha_2, \dots, \alpha_K)$, the aggregate property says that the first two vote shares and the sum of the remaining vote shares are distributed according to $\text{Dir}(\alpha_1, \alpha_2, \sum_{i=3}^{K}\alpha_i)$. Thus we have

$$\Pr(x_1 = x_2 = y) = \text{Dir}\big(y, y, 1 - 2y; sv_1, sv_2, s(1 - v_1 - v_2)\big). \tag{A.5}$$

Again using the aggregate property it can be shown that if we have K vote shares distributed

the three-candidate case using a Gaussian (though noting that a Dirichlet may be better). Hoffman (1982)'s Figure 1 is a beautiful representation of the integration necessary to compute pivot probabilities, again for the three-candidate case. Palfrey (1989) introduces a multinomial approach, later used by Herrmann, Munzert and Selb (2015).

[5] The Dirichlet model describes vote shares as continuous variables; multiplying by $1/n$ reconnects the model to the discrete character of elections.

according to Dir($s\mathbf{v}$), then

$$\Pr(x_3 = z | x_1 = x_2 = y) = \text{Beta}\left(\frac{z}{1-2y}; \alpha_3, \sum_{i=4}^{K} \alpha_i\right). \tag{A.6}$$

Note that the Beta distribution is just a special case of the Dirichlet distribution when $K = 2$. To calculate the probability that the vote share of party 3 is at z (given that parties 1 and 2 are both at y), we consider parties 3 to K to be dividing up the remaining $1 - 2y$ of vote share (such that we are evaluating the probability of party 3 at $z/(1-2y)$), and we use the aggregate property to lump together parties 4 through K.

Putting all of this together, we have that $\Pr(x_1 = x_2 = y, x_3 < y, \ldots, x_K < y)$ is approximately

$$\text{Dir}\big(y, y, 1-2y; sv_1, sv_2, s(1 - v_1 - v_2)\big) \prod_{i=3}^{K} \int_0^y \text{Beta}\big(\frac{z}{1-2y}; sv_i, s\sum_{j=3}^{K} v_j - sv_i\big) dz \tag{A.7}$$

To calculate the pivot probability for candidates 1 and 2, we substitute this expression into expression A.3 and numerically integrate across values of y between $1/K$ and $1/2$.

To demonstrate the validity of the method, we generate pivot probabilities for election results from UK constituencies in the 2005, 2010, and 2015 general elections at three levels of s using both our analytical/numerical approximation and a brute-force simulation method. For the simulation approach, we obtain 1 million draws for each election and we judge a near-tie for first to be a case where two parties have a vote share within one percentage point of each other and all other parties are lower; with this number of simulations (and non-optimized code), calculating pivot probabilities for a single value of s and a single election for 632 constituencies requires several hours. Our approach involves some calculation and numerical integration but is many hundreds of times faster. To validate our approach, we seek to establish that we can recover the pivot probabilities given by the simulation for cases where the simulation should work well (i.e., cases where the true pivot probabilities are not too small) and that our approach beats the simulation in other cases.

As a first validation, we compute the following for every election race using both the simulation approach and our approach:

$$\ln \frac{\pi_{Lab-Con}}{\pi_{LD-Con}} \tag{A.8}$$

where π_{A-B} is the probability of a tie for first between party A and party B. Figure A.1 compares this statistic as calculated by the simulation approach (horizontal axis) and our approach (vertical axis) for every constituency in the 2005, 2010, and 2015 general elections assuming $s = 8$ (left panel), $s = 20$ (center panel), and $s = 85$ (right panel). The results are clearly very similar in general.

Our method clearly outperforms a brute-force simulation approach in many cases not shown in Figure A.1: namely, those cases where the simulation method yielded a pivot probability of zero (and thus the log ratio of pivot probabilities is infinite or undefined). In these cases our approach is clearly preferred because it yields a positive pivot probability even for very unlikely events. To give a sense of the scale of the issue, consider the case where $s = 85$. Across the elections of 2005, 2010, and 2015, there are slightly more than 16,000 party pairings for which we can calculate a pivot probability; in the modal race there are five parties competing, meaning 10 unique pairings for which we can calculate a pivot probability. With our approach we obtain a positive pivot probability for all of these pairs. With the simulation approach and $s = 85$ we obtain a positive pivot probability for only around 4,500, or 28%; the remaining

Figure A.1: Our analytical/numerical approach recovers pivot probabilities produced by a simulation

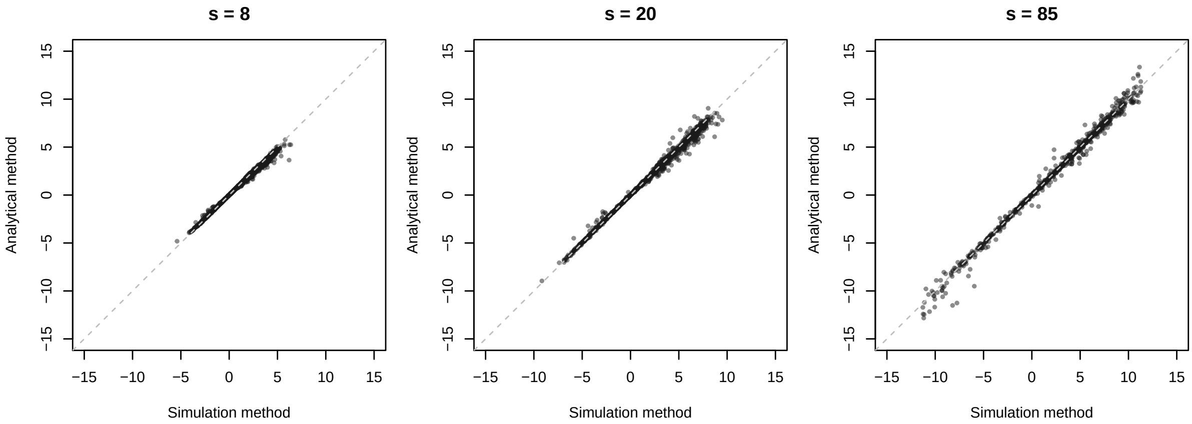

NOTE: We compute pivot probabilities for every pair of parties in each election in 2005, 2010, and 2015 using both a brute-force simulation approach and our analytical/numerical approach. In each panel, each dot shows the log of the ratio of the Lab-Con pivot probability to the LibDem-Con pivot probability for a single constituency contest using the simulation approach (horizontal axis) and our approach (vertical axis).

72% are zero. In many cases it may not matter whether the pivot probability is zero (as with the simulation approach) or a very small positive number (as with our approach), but it does matter when all of the pivot probabilities in a given case are quite small (as when one party has a comfortable lead). In such cases relative pivot probabilities produced by the simulation approach may depend heavily on random variation, because for a finite number of simulations the estimated probability could be zero or a very small number, which might have very different implications for analysis; also, the simulation method could yield zeros for all pivot probabilities in a given setting, in which case it is difficult to know how to proceed. Researchers may choose to ignore cases with very low absolute pivot probabilities (we do not), but they should not do so simply because the simulation method gives coarse estimates of very rare events.

A2: Formal relationship between τ and Myatt's Λ

Most prior research has used ad hoc approaches to identify the set of voters who might benefit from an insincere vote (e.g. those whose preferred candidate finishes third or lower) or to measure the intensity of the incentive to vote insincerely (e.g. using the electoral margin between the leading candidates or the gap in the like/dislike scores the voter assigns to the leading candidates). An important exception is the strategic incentive variable Λ introduced by Myatt (2000) and used in Fisher (2001), Herrmann, Munzert and Selb (2015), and Fisher and Myatt (2017). Fisher and Myatt (2017) defines Λ as $\frac{\pi_{23}-\pi_{13}}{\pi_{23}+2\pi_{12}}$, where e.g. π_{23} refers to the probability of a tie for first between the voter's second- and third-choice candidates; Fisher and Myatt (2017)'s Proposition 1 shows that the voter optimally votes for her second choice if and only if $\Lambda > \tilde{u} \equiv \frac{u_1-u_2}{u_1-u_3}$, where e.g. u_1 is the voter's utility from electing her first-choice candidate. Thus Λ encodes information about the voter's preference order and electoral context (but not preference intensity) into a single scalar measure which, when combined with another scalar measure $\tilde{u}$ characterizing relative preferences among the three candidates, indicates the optimal vote. Here we show that in the three-candidate case our measure of tactical incentives τ can be rewritten in terms of Myatt's Λ and $\tilde{u}$ as $\tau = (u_1 - u_3)(\pi_{23} + 2\pi_{12})(\Lambda - \tilde{u})$. Thus the sign of τ always indicates the same optimal vote that $\Lambda - \tilde{u}$ does, but τ additionally encodes information about the scale of preferences and pivot probabilities that is omitted in Myatt's approach.

We begin by writing out τ for the three-candidate case as

$$\begin{aligned}\tau &= \tilde{\mathbf{p}}(2) \cdot \mathbf{u} - \tilde{\mathbf{p}}(1) \cdot \mathbf{u} \\ &= \left[\pi_{13}u_1 + 2\left(\pi_{12} + \pi_{23}\right)u_2 + \pi_{13}u_3\right] - \left[2\left(\pi_{12} + \pi_{13}\right)u_1 + \pi_{23}u_2 + \pi_{23}u_3\right] \\ &= 2\pi_{12}(u_2 - u_1) + \pi_{23}(u_2 - u_3) - \pi_{13}(u_1 - u_3), \end{aligned} \tag{A.9}$$

adding and subtracting $\pi_{23}u_1$ to get

$$\tau = 2\pi_{12}(u_2 - u_1) + \pi_{23}(u_1 - u_3) - \pi_{23}(u_1 - u_2) - \pi_{13}(u_1 - u_3),$$

and rearranging as

$$\tau = (\pi_{23} - \pi_{13})(u_1 - u_3) - (\pi_{23} + 2\pi_{12})(u_1 - u_2).$$

Finally we divide both sides by $(u_1 - u_3)(\pi_{23} + 2\pi_{12})$, which produces

$$\tau = (u_1 - u_3)(\pi_{23} + 2\pi_{12})\left(\Lambda - \tilde{u}\right). \tag{A.10}$$

Because $(u_1 - u_3)$ is positive by definition, τ and $\Lambda - \tilde{u}$ must have the same sign: the two measures always agree about whether a tactical vote is optimal. Fisher and Myatt (2017) use Λ as a measure of the incentive to vote tactically for one's second choice. Inspecting equation A.10, τ can be seen to (i) incorporate information about relative preference ($\tilde{u}$) and (ii) preserve information about the scale of preferences ($u_1 - u_3$) and the absolute electoral stakes ($\pi_{23} + 2\pi_{12}$).

We see τ as distinct from Myatt's Λ in three principal ways. First, τ directly incorporates the electoral context and VNM preferences into a single scalar measure, whereas Myatt's approach deliberately separates the electoral context (summarized by Λ) from relative preferences ($\tilde{u}$). Whether this separation is desirable depends on the availability of a proxy for VNM utilities and one's model of voting behavior. Second, τ retains information about the absolute scale of tactical voting incentives (i.e. how much a voter stands to benefit or lose on the scale of the utility proxy by voting for a non-preferred candidate) that Myatt's approach discards; retaining scale information is desirable if one's model of voter behavior is the one in Section 2 above, but may not be in other cases. Third, Myatt's approach applies only to the three-candidate case, whereas τ can be computed for any number of candidates. In essentially three-candidate races where no adequate proxy for VNM preferences is available (such as the pre-2005 UK elections in Fisher and Myatt (2017)), Λ is the ideal scalar measure of tactical voting incentives; when a better preference measure is available (including in races with more than three candidates), τ captures in a single scalar measure both whether and to what extent a tactical vote is called for.

Appendix B: Testing explanations for heterogeneity in strategic responsiveness

What explains the heterogeneity in strategic responsiveness by social characteristics that is documented in the main text? Here we extend the basic regression specification reported above to provide some initial evidence concerning a number of different possible explanations. Each explanation we will consider involves a third *omitted variable* Z_i, which is itself associated with responsiveness to measured tactical incentives and is also correlated with the social characteristic of interest (or, alternatively, is differentially correlated with levels of τ for voters with and without the social characteristic of interest).

Our strategy for testing each explanation is to assess whether the interactions reported in Figure 5 of the main text – specifically, those estimated controlling for bins of τ – are attenuated when we re-estimate each regression model and add a control for Z_i and for the interaction between Z_i and $I(\tau_i > 0)$.

First, it may be that observed differences in responsiveness to tactical incentives across voters with and without a social characteristic of interest are really explained by variation in another of the five basic characteristics considered in this study. For example, the observed increase in strategic responsiveness by age may be driven by income, if higher income voters are more responsive to τ and if income is positively correlated with age.[6] Therefore, we examine how the estimated interaction between the social characteristic of interest and $I\{\tau_i > 0\}$ changes when we re-estimate the baseline model four times, each time controlling for one of the four remaining social variables and its interaction with $I\{\tau_i > 0\}$.

Second, we examine whether observed differences in strategic responsiveness by social characteristic are explained by variation in party support, either because supporters of different parties differ systematically in their objectives (expressiveness/farsightedness) or their perceptions of pivotal probabilities. We test for this by controlling for a series of indicators as to which party a voter most prefers and the interaction of these indicators with $I\{\tau_i > 0\}$.[7]

Third, we examine whether differences in strategic responsiveness by social characteristics may be driven by differences in intensity of party identification, a variable found to be strongly associated with likelihood of voting tactically in previous research (Niemi, Whitten and Franklin, 1992; Evans, 1994; Fisher, 2001). If party identification reflects an emotive attachment to a political party that is not fully reflected in like-dislike scores, voters with stronger party identifications may be less willing to vote tactically for a given value of τ. If, in addition, strength of party identifications is correlated with social characteristics, this may explain why voters with certain social characteristics are more or less responsive to tactically incentives. We test for this by controlling for strength of voters' self-reported party identification and its interaction with $I\{\tau_i > 0\}$.[8]

Fourth, we examine whether observed differences in strategic responsiveness by social characteristics arise because voters with certain social characteristics more accurately anticipate election outcomes in their constituencies. To assess this explanation, we control for an indicator

[6] It could also be, for example, that higher income voters are more strategically responsive and that income is positively associated with τ for old voters but negatively associated with τ for young voters, generating the observed association between age and responsiveness to τ.

[7] We set Conservative support as the reference category and include indicators for Labour, Liberal Democrat, UKIP, SNP, Plaid Cymru or Green support. For voters who indicate a clear most-preferred party in their like-dislike scores, we code this party as the one they support. Where there is no clear most-preferred party according to like-dislike scores, due to ties, but where a voter reports identifying with or feeling 'closer' to a party in response to party ID questions, we code party support based these latter questions.

[8] Strength of party identification is coded into four categories: no party ID, 'not very strong', 'fairly strong', 'very strong'. We set no party ID as the reference category in regressions.

measuring whether a respondent correctly anticipates which party will win their seat (and the interaction between this indicator and $I\{\tau_i > 0\}$). This is measured based on whether, when asked in the campaign wave of the BES how likely it was that each party would win the election in their constituency, a respondent assigns the highest likelihood to the party that ultimately won the seat.

Fifth, we examine whether observed heterogeneity in strategic responsiveness are explained by variation in campaign intensity. Party constituency campaigns parties try to mobilise tactical votes (Fisher, 2001) and it may be that voters with certain social characteristics tend to be located in areas where party election campaigns are more intense. We test this explanation by controlling for a number of alternative proxies for the election campaign intensity a respondent is likely to have experienced (and the interaction between each proxy and $I\{\tau_i > 0\}$): previous winning margin – the difference in vote share between the first and second-placed party in the respondent's constituency at the last election – which should be negatively related to campaign intensity; anticipated winning margin, according to contemporary poll-based forecasts; an indicator measuring whether a respondent reports being contacted by a political party in the past four weeks. We also subset the data to 2015 observations only and control for reported constituency campaign spending during, respectively, the long and short campaign.[9]

Sixth, we examine whether heterogeneity in strategic responsiveness by voter social characteristics is driven by variation in individuals' political knowledge, tendency toward instrumental decision-making, or perceived vote efficacy. We expect voters higher in each of these traits to display voting behavior that is more responsive to tactical incentives, and it could be that these traits are correlated with social characteristics.[10] Our measure of political knowledge is the proportion of correct answers a respondent gives to the domestic and international political knowledge batteries contained in the 2015 BES. Our measure of self-reported tendency toward instrumental political decision-making (hereafter labeled "strategic predisposition") is based on 2015 BES respondents self-reported level of agreement with two statements, "People should vote for the party they like the most, even if it's not likely to win" and "People who vote for small parties are throwing away their vote", recorded on a five-point scale. We take the average of a respondent's level of agreement to the two statements after reversing the polarity of response scale for the first statement. Our measure of respondent perceived vote efficacy is based on responses to the question, "How likely is it that your vote will make a difference in terms of which party wins the election in your local constituency?". The response scale was a 0-10 scale where 0 represents "very unlikely" and 10 represents "very likely".

Figure B.1 reports the results of this exercise. Each panel corresponds to a particular $I\{\tau_i > 0\}\times$ group membership interaction. In each panel, row 1 plots our 'baseline' estimate of this interaction, as well as the corresponding 95% confidence interval, based on regression equation 2 in the main text. Rows 2-12 show how the estimated interaction of interest changes when we re-estimate the baseline model, each time controlling for a different Z_i variable representing one of the possible explanations for heterogeneity in strategic responsiveness. In rows 14-18 (i.e. those highlighted in gray) we deal with explanations involving a Z_i variable that is only measured for the 2015 election data. Therefore, in row 13 we display a '2015 baseline' estimate of the $I\{\tau_i > 0\}\times$ group membership interaction, to serve as an appropriate point of comparison.[11]

The stability of estimates in rows 1-6 of each panel of Figure B.1 suggests that any observed

[9] Our measure is based on Electoral Commission measures of party campaign spending in each constituency as a percentage of the campaign spending limit for that constituency. We take the average score of the two top-spending parties in each constituency as our measure of spending intensity.

[10] It may also be that the correlation between these traits and observed levels of τ_i differs among voters with and without a social characteristic of interest, and that this drives observed variation in responsiveness to τ_i by social characteristic.

[11] These estimates are equivalent to the 2015 estimates displayed in Figure 7, main text.

Figure B.1: Sensitivity of estimated interactions to inclusion of controls

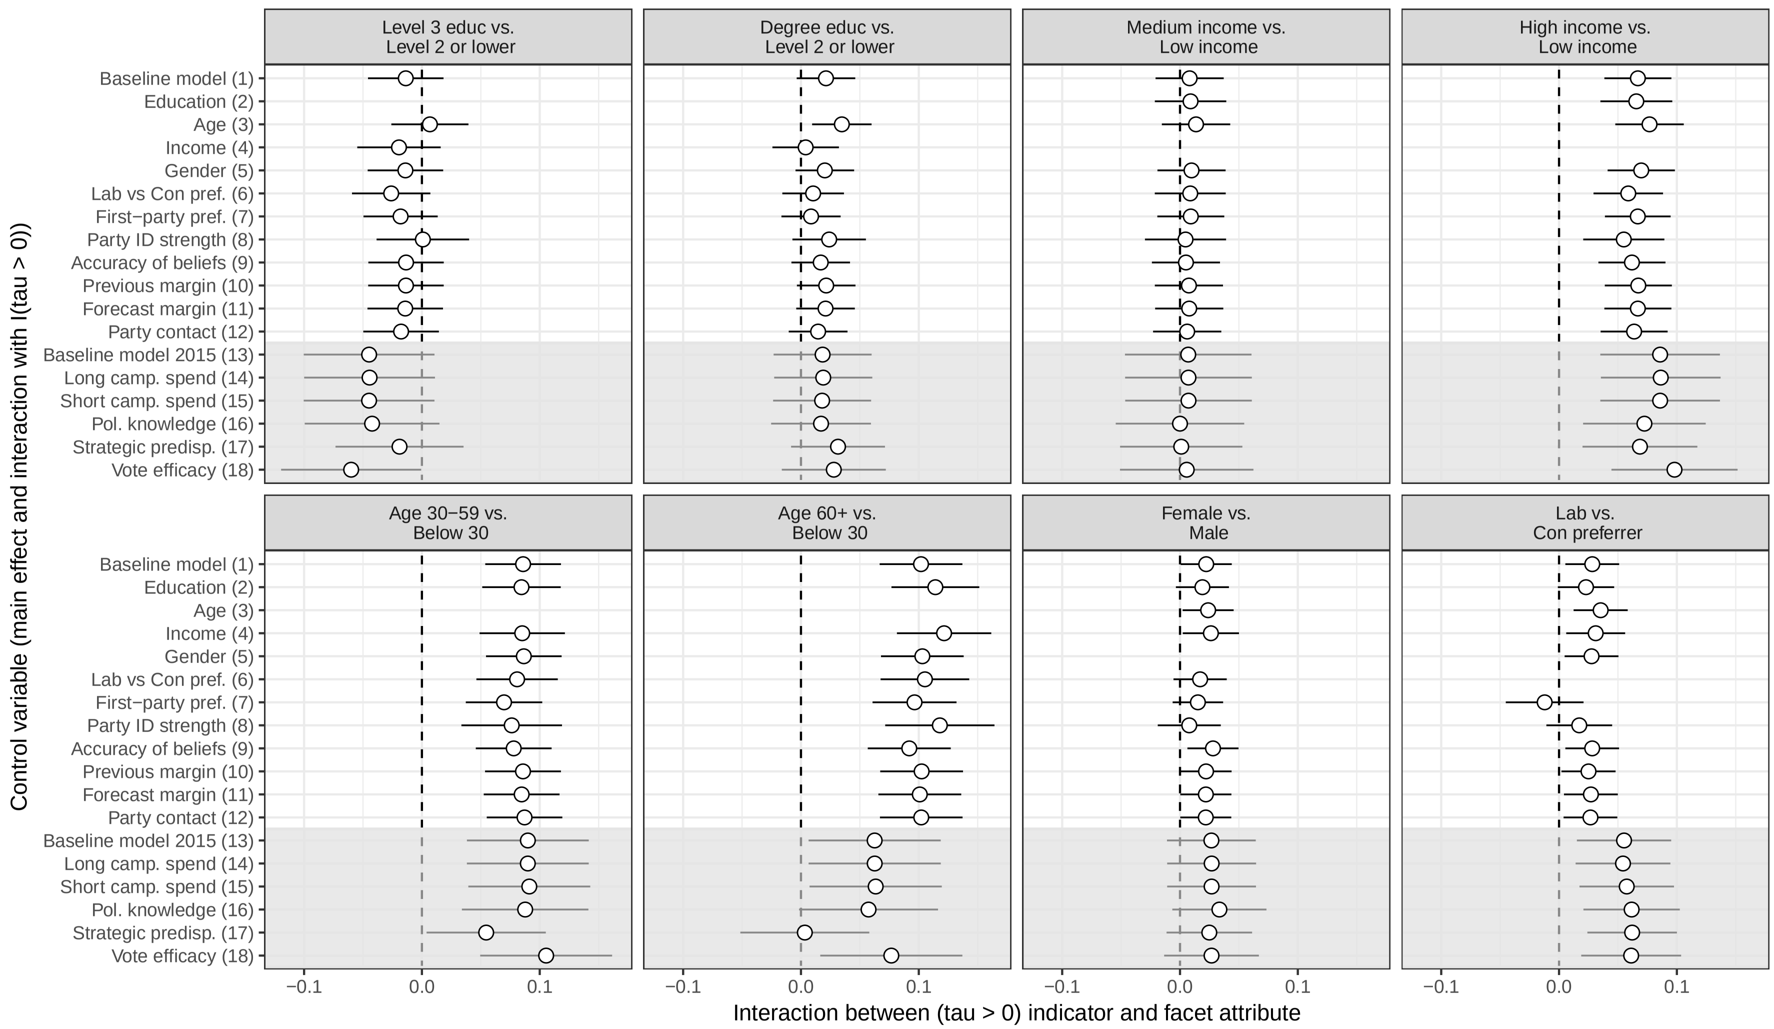

heterogeneity in strategic responsiveness by one of our five demographic or political variables of interest (education, income, age, gender and political leaning) is not well explained by variation in any of the other four remaining variables. Interestingly, however, inspection of row 3 indicates that, once we control for age and its interaction with $\tau > 0$, the estimated difference in strategic responsiveness between voters with the highest and lowest levels of education becomes positive and significant, though it remains relatively small.

Furthermore, in each panel, the point estimates of the interaction of interest change little from the relevant baseline estimate when we control for strength of party identification (rows 8 vs 1), accuracy of beliefs (row 9 vs 1), campaign intensity (rows 10-12 vs 1 and rows 14-15 vs 13), political knowledge (row 16 vs 13), or perceived vote efficacy (row 18 vs 13). Thus, observed heterogeneity in strategic behavior by social characteristic is not well explained by any of these factors.

However, comparison of rows 17 vs 13 in each panel does suggest that controlling for voters' self-reported strategic predisposition does somewhat attenuate some estimated interactions, particularly that between the age group indicators and $I\{\tau > 0\}$. This suggests that the increased strategic responsiveness of older voters that was detected in the main results may be attributable at least in part to older voters being more consciously instrumental in their vote decisions.

Finally, comparison of rows 7 and 1 in each panel shows that controlling for "first-party preferences" does induce a notable attenuation in the estimated interactions between the Labour Party preferrer indicator and $I\{\tau > 0\}$. Thus, the differences in strategic responsiveness between left- and right-leaning voters may be explained by associated differences in party support.

Do these various Z_i variables differ in their ability to explain observed heterogeneity in strategic responsiveness because some are themselves more or less strongly related to such responsiveness? To answer this question, we now turn to report the results of a series of regressions where we model best insincere voting as a function of each Z_i variable and its interaction with $I\{\tau > 0\}$, dropping social characteristics from the model specification. Specifically, we estimate the regression equation

$$E[Y_i] = g(\tau_i, \text{Year}_i) + \beta_1 Z_i + \beta_2 I\{\tau_i > 0\} + \beta_3 Z_i \times I\{\tau_i > 0\}. \qquad \text{(A.11)}$$

The control function $g(\tau_i, \text{Year}_i)$ includes indicators for deciles of τ in the British electorate and – in models which pool observations across elections – indicators for election years and their interaction with τ bins. The main coefficient of interest in Equation A.11 is β_3, which measures the change in responsiveness to $I\{\tau_i > 0\}$ when Z_i increases by one unit.

Table B.1 shows coefficient estimates when the Z_i variables are first party preference (column 1) and strength of party identification (2). The interaction terms in column 1 show that voters who most-prefer Labour, UKIP and the Greens are significantly more responsive to tactical incentives than are those voters who most prefer the Conservative Party. The interaction terms in column 2 indicate that voters with strong party identification are less responsive to tactical incentives in their voting behaviour than are voters who do not identify with any party. This is broadly consistent with the notion that party indentification reduces strategic behaviour.

Table B.2 shows coefficient estimates for other Z_i variables discussed in the main text. Column 1 shows that voters with more accurate beliefs about the election outcome in their seat are, as expected, significantly more responsive to tactical incentives.

Turning to proxies for campaign intensity, Column 2 shows that voters in seats that were less marginal in the previous election are not significantly more or less strategically responsivene, whereas Column 3 shows that voters in seats *forecast* to be less marginal are marginally more strategically responsive. On the one hand this result is puzzling given we would expect such

Table B.1: Heterogeneity in strategic responsiveness by respondent party support and strength of party identification

	(1)	(2)
1st pref Lab	0.011**	
	(0.003)	
1st pref LD	0.005	
	(0.005)	
1st pref SNP	−0.009	
	(0.005)	
1st pref PC	0.076*	
	(0.038)	
1st pref UKIP	−0.012	
	(0.010)	
1st pref Grn	0.002	
	(0.033)	
PID weak		0.025*
		(0.010)
PID moderate		−0.010
		(0.010)
PID strong		−0.035***
		(0.010)
$I(\tau > 0)\times$ 1st pref Lab	0.137***	
	(0.020)	
$I(\tau > 0)\times$ 1st pref LD	0.006	
	(0.018)	
$I(\tau > 0)\times$ 1st pref SNP	0.041	
	(0.043)	
$I(\tau > 0)\times$ 1st pref PC	0.047	
	(0.061)	
$I(\tau > 0)\times$ 1st pref UKIP	0.178***	
	(0.021)	
$I(\tau > 0)\times$ 1st pref Grn	0.294***	
	(0.037)	
$I(\tau > 0)\times$ PID weak		−0.020
		(0.029)
$I(\tau > 0)\times$ PID moderate		−0.015
		(0.029)
$I(\tau > 0)\times$ PID strong		−0.121***
		(0.032)
Constant	0.028***	0.047***
	(0.008)	(0.012)
Control for (binned) τ?	Yes	Yes
Control for election year?	Yes	Yes
Control for (binned) τ × election year?	Yes	Yes
Observations	24,922	16,224
R^2	0.259	0.227
Adjusted R^2	0.257	0.226
Note:	*p<0.05; **p<0.01; ***p<0.001	

voters to receive less intensive campaigns. On the other hand, it may be driven by a process whereby forecast winning margin itself captures intensity of tactical incentives for a given τ decile. Column 4 shows that voters who report having been contacted by a party during the election campaign are more responsive to τ. Columns 5 and 6 show that neither local long campaign spending nor local short campaign spending are significantly associated with strategic responsiveness.

Turning to voter attributes, column 7 shows that voters who score higher in the BES political knowledge test are not significantly more responsive to τ than voters who score lower. In line with expectations, however, column 8 shows that a voter's self-reported strategic predisposition is strongly positively associated with strategic responsiveness, while column 9 shows that voters who have a greater sense of vote efficacy are also more strategically responsive.

Table B.2: Heterogeneity in strategic responsiveness by additional voter and constituency attributes

	Belief accuracy	Prev. margin	Fcast margin	Party contact	Long spend	Short spend	Knowledge	Strat. disp.	Efficacy
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Main effect	−0.009**	−0.001***	−0.001***	−0.010**	0.0001	−0.00002	−0.001	−0.011**	−0.020***
	(0.003)	(0.0002)	(0.0002)	(0.003)	(0.0001)	(0.00002)	(0.005)	(0.004)	(0.005)
Interaction with $I(\tau > 0)$	0.073***	0.001	0.001*	0.068***	−0.0002	−0.0001	0.032	0.326***	0.096***
	(0.012)	(0.001)	(0.001)	(0.012)	(0.0004)	(0.0001)	(0.021)	(0.018)	(0.022)
Constant	0.039***	0.041***	0.040***	0.038***	0.014*	0.023**	0.020***	0.021***	0.023***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.007)	(0.007)	(0.005)	(0.005)	(0.006)
Years	All	All	All	All	2015	2015	2015	2015	2015
Control for (binned) τ?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Control for election year?	Yes	Yes	Yes	Yes	No	No	No	No	No
Control for (binned) $\tau \times$ election year?	Yes	Yes	Yes	Yes	No	No	No	No	No
Observations	24,922	24,910	24,922	24,922	8,606	8,606	8,260	8,333	7,309
R^2	0.231	0.228	0.228	0.230	0.246	0.247	0.245	0.306	0.246
Adjusted R^2	0.230	0.227	0.227	0.229	0.245	0.246	0.244	0.305	0.245

Note: $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

Appendix C: Robustness checks

Exploring consequences of missing data

We note in the main text that, for our main analysis, we drop BES respondents for whom we are unable to construct a measure of tactical incentive, τ. Could our main results concerning differences in strategic responsivness (SR) by age and income be driven by patterns of missingness on τ? This could occur if missingness on τ is correlated with both social characteristics and SR or if missingness on τ is differentially correlated with SR depending on social charateristics. Here we examine the extent to which these two types of processes could explain our main findings.

Before beginning the analysis let us be clear as to the three possible reasons why a respondent might receive a missing τ score and therefore be dropped from our analysis. First, approximately half of 2015 BES respondents were not asked the party rating battery we use to measure outcome-based utility (with the other half randomly allocated propensity-to-vote questions instead) and so have missing τ scores. We ignore these respondents here because receipt of the party ratings battery was randomized, such that this type of missingness is in expectation uncorrelated with voter characteristics. Second, of the 29,226 respondents in the pooled BES sample who received the party rating battery, 985 (3%) provided a party rating for an insufficient number of parties (less than three) and consequently have a missing τ score. Third, 1,455 (5%) were deemed to have inconsistent preferences - because the respondent indicates two or more parties as tied for first place in their party ratings but then says they identify with or are closer to an entirely different party when responding to questions we use to break the tie - and consequently have a missing τ score. Note that respondents could offer both an insufficient number of party ratings and exhibit inconsistent preferences; indeed 72 did so.

Figure C.1: **Differences in rate of missing scores on τ by social characteristic**

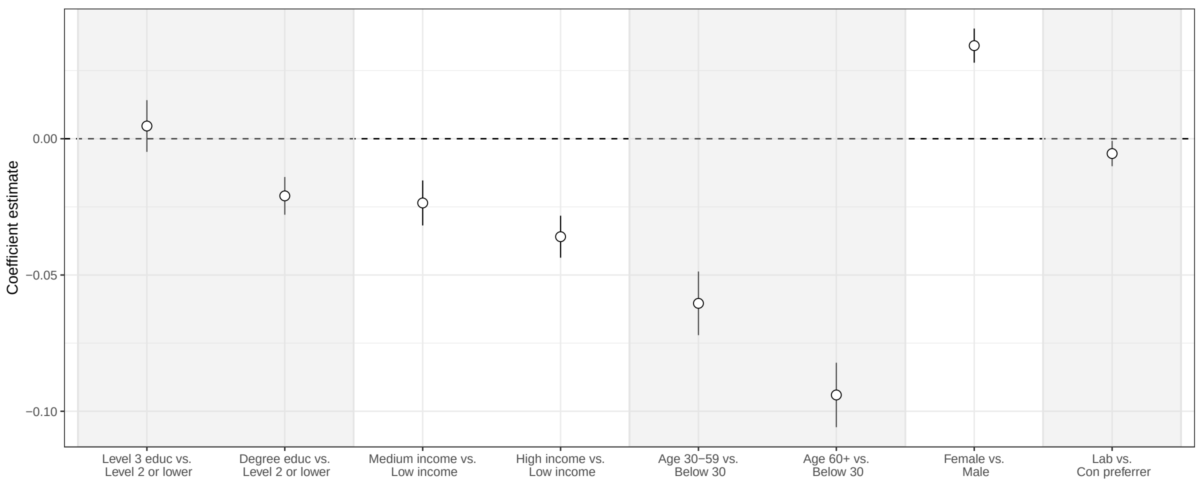

Figure C.1 shows results when we construct an indicator of missingness on τ due to either insufficient party ratings or inconsistent preferences and regress this on respondent social characteristics (estimating a seperate model for each social characteristic). Clearly, missingness on τ is associated with several social characteristics. For example, the rate of missingness on τ is 3 points lower for high income respondents than for low income respondents, and is around 6 and 9 points lower for middle and old age respondents, respectively, compared to young respondents.

Given that, in our sample, lower income and younger respondents are more likely to have

Figure C.2: **Strategic response function by number of party ratings questions answered by respondent**

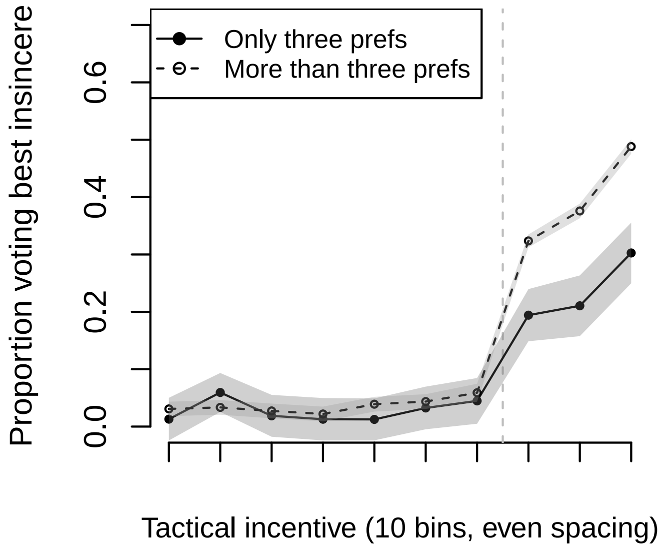

missing τ scores, one way in which dropping due to missingness on τ could generate our main findings (that younger and lower income voters are also less responsive to strategic incentives) is if true SR is positively associated with missing τ scores. For example, if true SR and being young are both positively associated with missingness on τ, analysis of respondents with non-missing τ scores would underestimate SR for young voter relative to SR for older voters.

However, we find it hard to see why survey item non-response (answering an insufficient number of party ratings) and inconsistent preferences – the two causes of missing τ values considered here – would be associated with greater strategic awareness and/or behavior. Furthermore, while we cannot directly estimate SR for respondents with missing τ scores and compare this to (observed) SR among respondents with non-missing τ scores, we can provide other evidence suggesting that respondents with missing τ scores are on average no more – and probably less – strategic than other respondents. First, Figure C.2 shows how, among respondents with non-missing τ scores, respondents with only three non-missing party ratings (who were essentially on the verge of being dropped) were substantially less responsive to strategic incentives than respondents with four or more non-missing party ratings (who were not on the verge of being dropped). It seems unlikely that those with missing τ scores are more strategic than those with non-missing scores when, among the latter, those on the verge of being assigned a missing τ score are *less* strategic than those who are not. Second, Table C.1 presents regression results for 2015 showing how missingness on τ is associated with the three respondent attributes that we find to be positively associated with SR in Appendix B, Table B.2: accuracy of beliefs about the election outcome in their constituency, self-reported strategic predisposition, and self-reported sense of vote efficacy. It shows that respondents with missing τ scores have significantly less accurate beliefs, significantly lower vote efficacy, and no greater strategic disposition. It thus appears that, if anything, missingness on τ is likely to be associated with *lower* levels of SR. If this is the case, then higher rates of missingness among younger and lower income voters cannot alone account for the fact that we find these voters to be *less* strategically

Table C.1: Association between missingness on τ and variables shown to be predictive of strategic responsiveness in Appendix B (2015 BES respondents)

	Winner correct	Strat. disposition	Efficacy
	(1)	(2)	(3)
Missing τ	-0.201^{***}	-0.019	-0.136^{***}
	(0.017)	(0.017)	(0.021)
Constant	0.627^{***}	0.004	0.008
	(0.005)	(0.005)	(0.006)
Observations	10,140	9,539	8,280
R^2	0.015	0.0001	0.005
Adjusted R^2	0.015	-0.00000	0.005
Note:		*p<0.05; **p<0.01; ***p<0.001	

responsive in our main analysis.

Might our results instead be explained by an alternative pattern of missingness on τ, where younger and lower income voters appear less strategic because of differential associations between missingness on τ and true SR by age and income groups? For example, if the association between missingness on τ and true SR is positive among young people but negative among old people, analysis of respondents with non-missing τ scores would underestimate young voters' SR relative to that of older voters.

To examine the sensitivity of our age and income results to this second type of missingness process, we perform data imputation for respondents with missing scores on τ. We create multiple versions of our data, each with imputation for missing scores performed according to a different assumption – captured by a parameter we label c – about how the relationship between SR and missingness on τ varies with the social characteristic of interest. Each version of the (partially) imputed data allows us to re-estimate how raw SR varies by the social characteristic of interest when c takes on a particular value.[12]

More specifically, let there be two groups on the social characteristic of interest: the *High* group is that estimated to have higher relative SR in our main analysis, while the *Low* group is that estimated to have lower SR in our main analysis. Furthermore, let M_i be a binary indicator equal to one when τ_i is missing for respondent i and zero otherwise. We define a parameter $c \geq 0$, which captures the assumed difference in average SR for respondents in the *Low* group for whom $M_i = 1$ and respondents in the *Low* group for whom $M_i = 0$. Furthermore, for the *High* group, we set the difference in average SR between respondents with $M_i = 1$ and $M_i = 0$ to be $-c$. Together these assumptions imply that the difference in SR between respondents with and without missing τ in the *Low* group is nonnegative, that the difference in SR between respondents with and without missing τ in the *High* group is nonpositive, and that $2c$ is the difference-in-difference in SR – i.e., the shift in the relative SR of the *Low* and *High* groups when moving from the sub-sample with non-missing τ to the sample with missing τ. As $2c$ increases, our imputation process assumes a greater reversal in the relationship between SR and *High* vs *Low* group membership when $M_i = 1$ compared to that observed in the sub-sample for whom $M_i = 0$.

Our imputation process for respondents with missing τ scores conditions on c, survey year (t) and grouping on the social characteristic of interest (*High* vs *Low*). It assumes that, conditional on year, the distribution of $I\{\tau > 0\}$ – the positive tactical incentive indicator – is the same for

[12] We estimate differences in raw SR across groups by using OLS to estimate equation 2 in the main text, leaving out the $g(\tau_i)$ term.

respondents with missing τ scores as for respondents with observed τ scores. It also assumes that, when their imputed $I\{\tau > 0\} = 0$, the rate of best non-sincere voting (y) for $M_i = 1$ respondents from a given year and social group is the same as the observed rate among $M_i = 0$ respondents from the same year and social group and for whom $I\{\tau > 0\} = 0$. When imputed $I\{\tau > 0\} = 1$, the imputation process again uses the observed rate of best insincere voting among comparable $M_i = 0$ respondents (i.e., from the same year and social group and for whom $I\{\tau > 0\} = 1$) as the basis for the assumed rate of best insincere voting among $M_i = 1$ respondents. However, it adds or subtracts c to this base rate depending on whether the social group is *Low* or *High*, before performing imputation.

The specific imputation steps are as follows:

1. Subset to the sub-sample of $N_{M=1}$ respondents for whom $M_i = 1$
2. For each year t:
   (a) Subset to the sub-sample of $N_{M=1,t}$ respondents from t
   (b) **Assignment of $I\{\tau_i > 0\}$ values:** Assign $I\{\tau_i > 0\} = 1$ to $100\phi_{M=0,t}$% of respondents in the sub-sample, where $\phi_{M=0,t}$ is the observed mean of $I\{\tau > 0\}$ among respondents with $M_i = 0$ in year t. Assign $I\{\tau_i > 0\} = 0$ to remaining respondents in the sub-sample.
   (c) **Assignment of the best insincere vote indicator, y:**
      i. Among those for whom $I\{\tau_i > 0\} = 0$ and who are in the *Low* group, assign $y_i = 1$ to $100\psi_{M=0,t,\text{Low},\tau\leq 0}$% of respondents, where $\psi_{M=0,t,\text{Low},\tau\leq 0}$ is the observed mean of y among respondents in the *Low* group with $M_i = 0$ in year t for whom $\tau_i \leq 0$. Assign $y_i = 0$ to the remainder.
      ii. Among respondents for whom $I\{\tau_i > 0\} = 0$ and who are in the *High* group, assign $y_i = 1$ to $100\psi_{M=0,t,High,\tau\leq 0}$% of respondents, where $\psi_{M=0,t,High,\tau\leq 0}$ is the observed mean of y among respondents in the *High* group with $M_i = 0$ in year t for whom $\tau_i \leq 0$. Assign $y_i = 0$ to the remainder.
      iii. Among respondents for whom $I\{\tau_i > 0\} = 1$ and who are in the *Low* group, assign $y_i = 1$ to $100 \times (\psi_{M=0,t,Low,\tau>0} + c)$% of respondents, where $\psi_{M=0,t,Low,\tau>0}$ is the observed mean of y among respondents in the *Low* group with $M_i = 0$ in year t for whom $\tau > 0$. Assign $y_i = 0$ to the remainder.
      iv. Among respondents for whom $I\{\tau_i > 0\} = 1$ and who are in the *High* group, assign $y_i = 1$ to $100 \times (\psi_{M=0,t,High,\tau>0} - c)$% of respondents, where $\psi_{M=0,t,High,\tau>0}$ is the observed mean of y among respondents in the *High* group with $M_i = 0$ in year t for whom $\tau > 0$. Assign $y_i = 0$ to the remainder.
3. Recombine the year-sepcific sub-samples of respondents with $M_i = 1$ together with the sub-sample of respondents for whom $M_i = 0$.
4. Estimate difference in raw SR between the *High* and *Low* groups for the resulting sample.

Figure C.3 summarizes key results from running this imputation process with respect to the social characteristic age, treating the 'Below 30' age-group as the *Low* group and the '30-59' and '60+' age-groups as the *High* groups. The x-axis shows values of $2c$ between 0 and 0.4, such that points further to the right correspond to data where missing scores are imputed assuming a greater reversal in the relationship between SR and age group when τ_i is missing compared to that observed when τ_i is non-missing. The height of the black line measures the estimated difference in SR comparing the 30-59 age group to the below 30 age group (top panel) or the 60+

Figure C.3: **How does estimated heterogeneity in SR by age group change when we impute missing data based on different assumptions?**

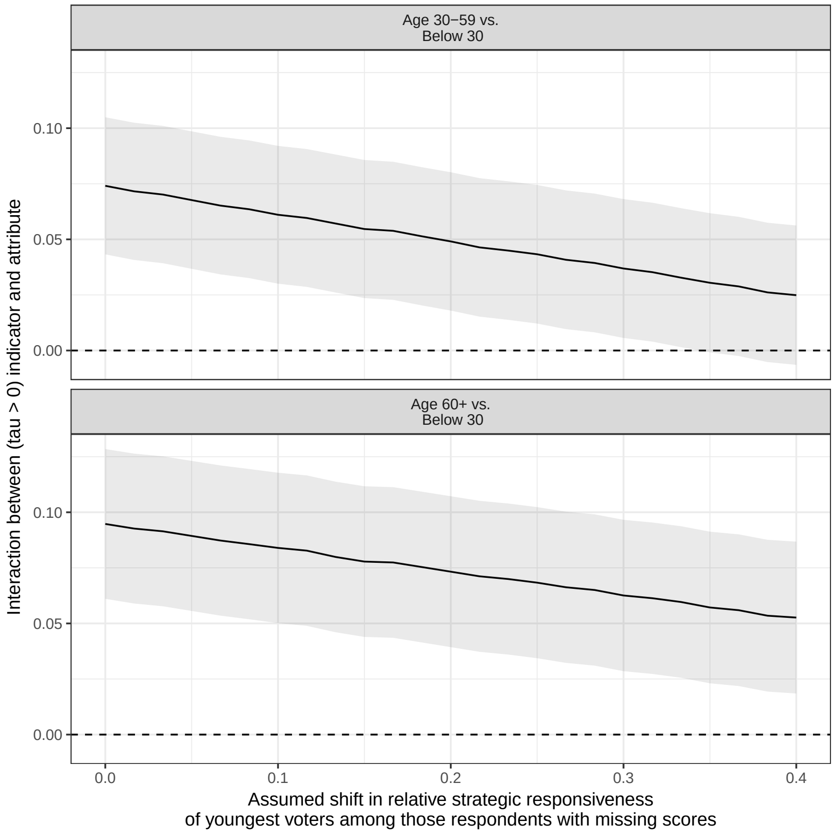

age group to the below 30 age group (bottom panel). Gray bands represent 95% confidence intervals. Recall that in the equivalent main analysis (open circles in Figure 5 in the main text) SR was estimated to be around 7.5 and 9 points higher for voters aged 30-59 and 60+, respectively, than for voters aged below 30, and both differences were significant. Figure C.3 shows that, as $2c$ increases, these differences decline, but remain positive and significant until $2c$ becomes quite extreme. The estimated difference in SR comparing voters aged 60+ to those below 30 (top panel) is still positive and significant when $2c = .4$. The estimated difference in SR comparing voters aged 30-59 to those below 30 (top panel) only becomes non-significant when $2c > .33$. For $2c > .33$ to hold, the difference in SR for young respondents with missing τ and young respondents with non-missing τ would have to be at least 0.16 points and the assumed difference in SR for middle-aged respondents with missing τ and middle-age respondents with non-missing τ would have to be -0.16 points or lower. Both of those differences in SR are larger in absolute magnitude than any differences in SR we observe by social characteristics in our main analysis. We therefore regard their joint occurrence highly unlikely. Consequently, we conclude that our main findings concerning heterogeneity in SR by age group are unlikely to be driven by patterns of missingness on τ.

Figure C.4 summarizes key results from running the imputation process with respect to the social characteristic income, treating the 'Low income' and 'Medium income' groups as the *Low* group and the 'High income' groups as the *High* group. In the equivalent main analysis (open circles in Figure 5 in the main text) estimated SR was non-significantly different for low

Figure C.4: **How does estimated heterogeneity in SR by income group change when we impute missing data based on different assumptions?**

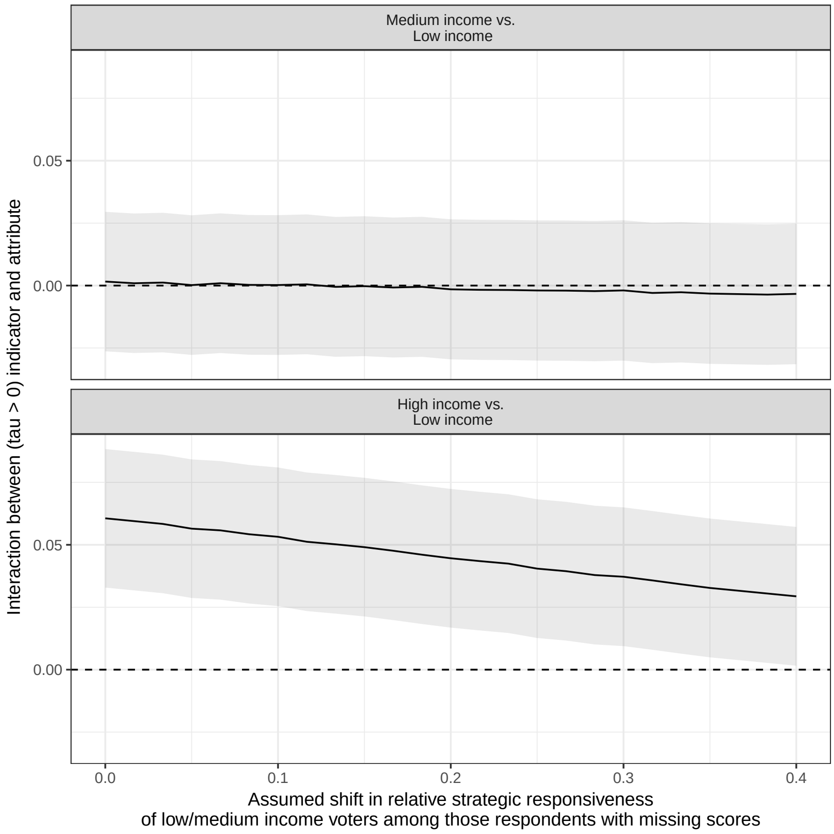

and medium income voters (with a point estimate of around zero for the difference), but was estimated to be significantly different – and around 6 points higher – for high income voters compared to low income voters. Figure C.3 shows that, as $2c$ increases, the estimated difference in SR between medium and low income voters (top panel) stays around zero and non-significant. The estimated difference in SR between high and low income voters (bottom panel) declines, but remains positive and significant even when $2c = 0.4$. We therefore conclude that our main findings concerning heterogeneity in SR by income are again unlikely to be driven by patterns of missingness on τ.

Differential measurement error

Readers may wonder whether the differences in strategic responsivness between social groups reported in the main text arise because of differential measurement error. For example, do young people appear less strategic than old people because our measure of younger voters' utility is noisier than that for older voters? This would lead to more measurement error in τ for younger voters, and lower estimates of strategic responsiveness even if younger and older voters were in fact equally strategic. Here we discuss several analyses which test plausible ways in which such differential might arise and which suggest that it does not drive the main differences in strategical responsivness that we report in the main text.

First, one could imagine that different social groups may place more or less weight on party leaders versus parties in general. For example, young voters' utility from an electoral outcome may depend relatively more on what they think of party leaders, which is an aspect of preferences we ignored in our main analysis. If that were the case, then our utility measure would be a worse proxy for young voters' preferences than it is for older voters' preferences, which could explain the difference in strategic responsiveness we find. (The same argument of course applies to e.g., poorer and richer voters.) To check this, we reproduce the analysis using leader ratings rather than party ratings as our proxies for utility, focusing on 2015, the only year in which respondents were asked to rate all leaders in the post-election survey. Figure C.5 shows the results of this exercise. Whether we use party leader ratings (empty diamonds) or the utility measure based on party ratings from the main text (filled gray circles), key differences in strategic responsiveness persist: younger voters and poorer voters appear to be less strategic. The difference between Labour and Conservative preferrrers does becomes non-significant, however, when using party leader ratings.

Figure C.5: **Sensitivity of key results to different utility measures (2015 only)**

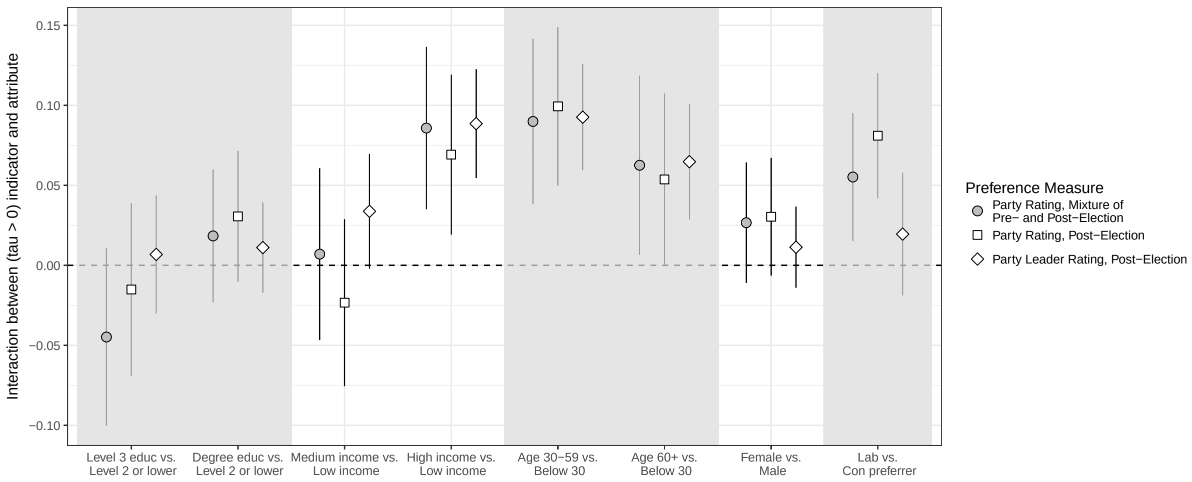

Second, one could imagine that the opinions of individuals in certain social groups are more changeable during the course of the campaign than are those of individuals in other social groups. If, for example, young voters' opinions are more changeable, pre-election party ratings (which we use for the Greens and UKIP, because post-election ratings were not asked in 2005 and 2010) are a noisier proxy of preference for young voters' preferences than for older voters' preferences. This would lead to a difference in measured strategic responsiveness even if the two groups were equally strategic. To check this, in Figure C.5 we also reproduce the 2015 analysis (filled gray circles) using only post-election party ratings (empty squares). Once more, we find that key differences in strategic responsiveness found in the main analysis persist. Voters age 30-59 are again significantly more strategically responsive than voters aged below 30. The difference in strategic responsiveness of voters aged 60 or above and those aged below 30 becomes marginally non-significant when using only post-election party ratings, but the point estimate is only slightly attenuated. High income voters remain significantly more strategically responsive than low income voters.

Third, one could imagine that inviduals with certain social characteristics may pay more attention to local candidates than the party as a whole. This would again lead to differential measurement error in tau and differences in measured strategic responsiveness, with groups

Figure C.6: **Differences in rate at which local candidate is mentioned as reason for vote choice**

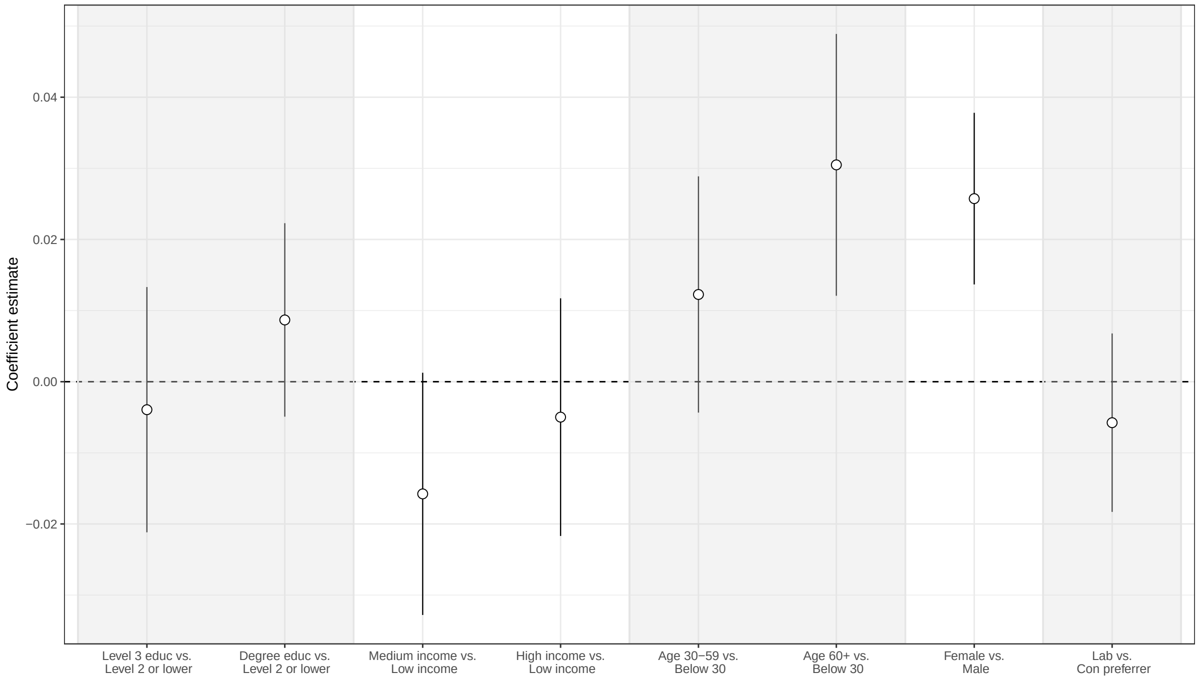

who pay more attention to candidates expected to appear less strategically responsive. To check this, we create an indicator variable measuring whether respondents who, when asked to explain why they chose the party they did, chose a response option indicating that the party had the "best candidate". Figure C.6 shows results when we regress this indicator on our social characteristic indicators (with seperate models for each social characteristic). Contrary to this potential explanation for the observed differences in strategic responsivness in the main text, there is no strong indication that the groups found to be least strategically responsive in our main results were more likley to vote on the basis of local candidates. Voters over 60 are significantly more likely than those under 30 to vote on the basis of local candidates, but were also found to be more strategically responsive in the main analysis. Females are also significantly more likely than males to vote on the basis of local candidates, but also tended to have higher estimated strategically responsiveness in the main analysis (although the difference in strategic responsiveness of females and males varied in significance across specifications). Remaining social group indicators are not significantly associated with voting on the bases of local candidates.

Fourth, elements of the analysis presented in Appendix B – where we looked at potential substantive explanations for observed heterogeneity in strategic responsiveness – speak to the differential measurement error concern. In particular, in Figure B.1 we tested whether our main estimated differences in strategic responsiveness between social groups (row 1) are attenuated when controlling for the main effect of either the party that the voter supports (row 7) or the strength of a voter's party identification as well as their interaction with a $I\{\tau_i > 0\}$ indicator. It is plausble that either of these variables may be associated with differential measurement error in the mapping of preferences to like-dislike scores: supporters of some parties, or respondents with stronger party ID, may tend to systematically over- or understate the differences in the utility they receive from their most-preferred and second-best party winning their seat. The results in Figure B.1 showed that controlling for these variables did not substantially attenuate differences in strategic heterogeneity by age or income, but that controlling for party support did somewhat attenuate estimated differences in strategic responsiveness comparing Labour and Conservative Party preferrers. Thus, while estimated differences in strategic responsiveness between left- and right-leaning voters may potentially be explained by differential measurement error associated with party support, other estimated differences by age and income do not.

Results are similar when we use different outcome variables

The top panel of Figure C.7 reports the same sensitivity analysis as Figure 7 in the main text, which is based on analysis where the outcome variable is whether a voter casts a best incenere vote or not. The second and third panels of Figure C.7 report the same analysis but with different measures of voting behaviour as the outcome variable. In the second panel the outcome variable is voting for *any* party other than one's most preferred party. The results are essentially identical to those in the top panel. In the third panel of Figure C.7 the outcome variable is Fisher's (2004) measure of tactical voting (which is based on respondent self-reported reasons for vote choice). With this alternative outcome variable, strategic responsiveness continues to vary significantly and substantially by voter age. The differences in strategic responsiveness by income and political leaning are in the same direction as in the main results, although the pooled estimates of these differences are now marginally non-significant at the 0.05 level. Differences are smaller in magnitude as would be expected given the lower rate of measured tactical voting.

Figure C.8 shows how the proportion of voters casting any insincere vote varies with τ. This proportion is higher than the proportion casting the best insincere vote (Figure 6 in the main text) but shows the same monotonic relationship and the same differences by social characteristic.

Results are similar when our model of counterfactual election outcomes is based on forecasts rather than observed results

In the analysis in the paper and in the sensitivity analysis shown in 7 (main text), we calculated τ based on a model of counterfactual elections whose expected outcome is the actual outcome. Alternatively, we could use forecasted results as the basis of our model of counterfactual election outcomes. The bottom panel of Figure C.9 reports the same sensitivity analysis as in the main text (replicated in the top panel of the Figure), but using forecasts as the source of expected election results. Concentrating on the pooled estimates, as in the main analysis, the youngest group of voters have significantly lower levels of strategic responsiveness than older voters, high income voters have significantly higher levels of strategic responsiveness than low income voters, and left-leaning voters have significantly higher levels of strategic responsiveness than right-leaning voters. Interestingly, in this analysis, we also find evidence that females have a significantly higher level of strategic responsiveness than males.

Results are similar when we increase the uncertainty in our model of counterfactual elections

In the analysis in the paper and in the sensitivity analysis above, we calibrated the precision of our model of counterfactual elections by setting the mean at the observed result and choosing the Dirichlet precision parameter s to maximize the likelihood of election forecasts; this led to a choice of $s = 85$. The bottom panel of Figure C.10 reports the same sensitivity analysis as Figure 7 (main text), but in the other two panels we repeat the analysis for two alternative values of the

Figure C.7: **Sensitivity analysis by alternative outcome variables**

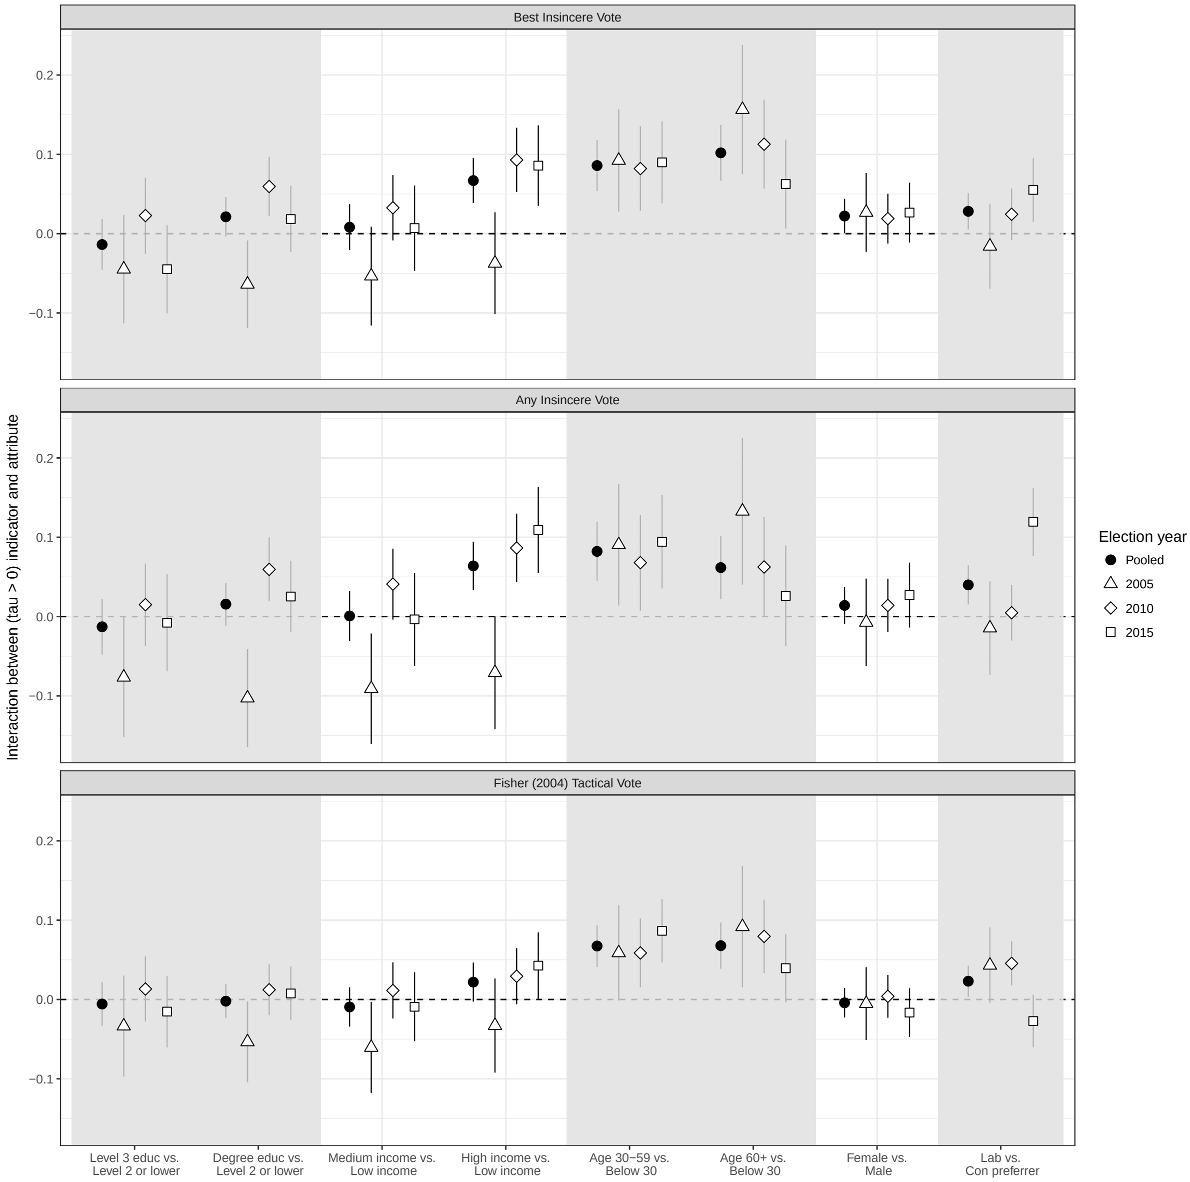

precision parameter, $s = 12$ and $s = 20$. $s = 12$ is of interest because it is the estimate of actual precision in the British electorate arrived at in Fisher and Myatt (2017). With $s = 20$, the standard deviation of each party's vote share is roughly double that at $s = 85$. Concentrating on the pooled estimates, across all values of s shown, younger voters are estimated to have significantly lower levels of strategic responsiveness than older voters and high income voters are estimated to have significantly higher levels of strategic responsiveness than low income voters. Regarding political leaning, left-leaning voters are consistently estimated to have higher levels of strategic responsiveness than right-leaning voters across all levels of s shown, although the difference is marginally non-significant when $s = 20$.

Figure C.8: Strategic response functions with “any insincere vote” as outcome

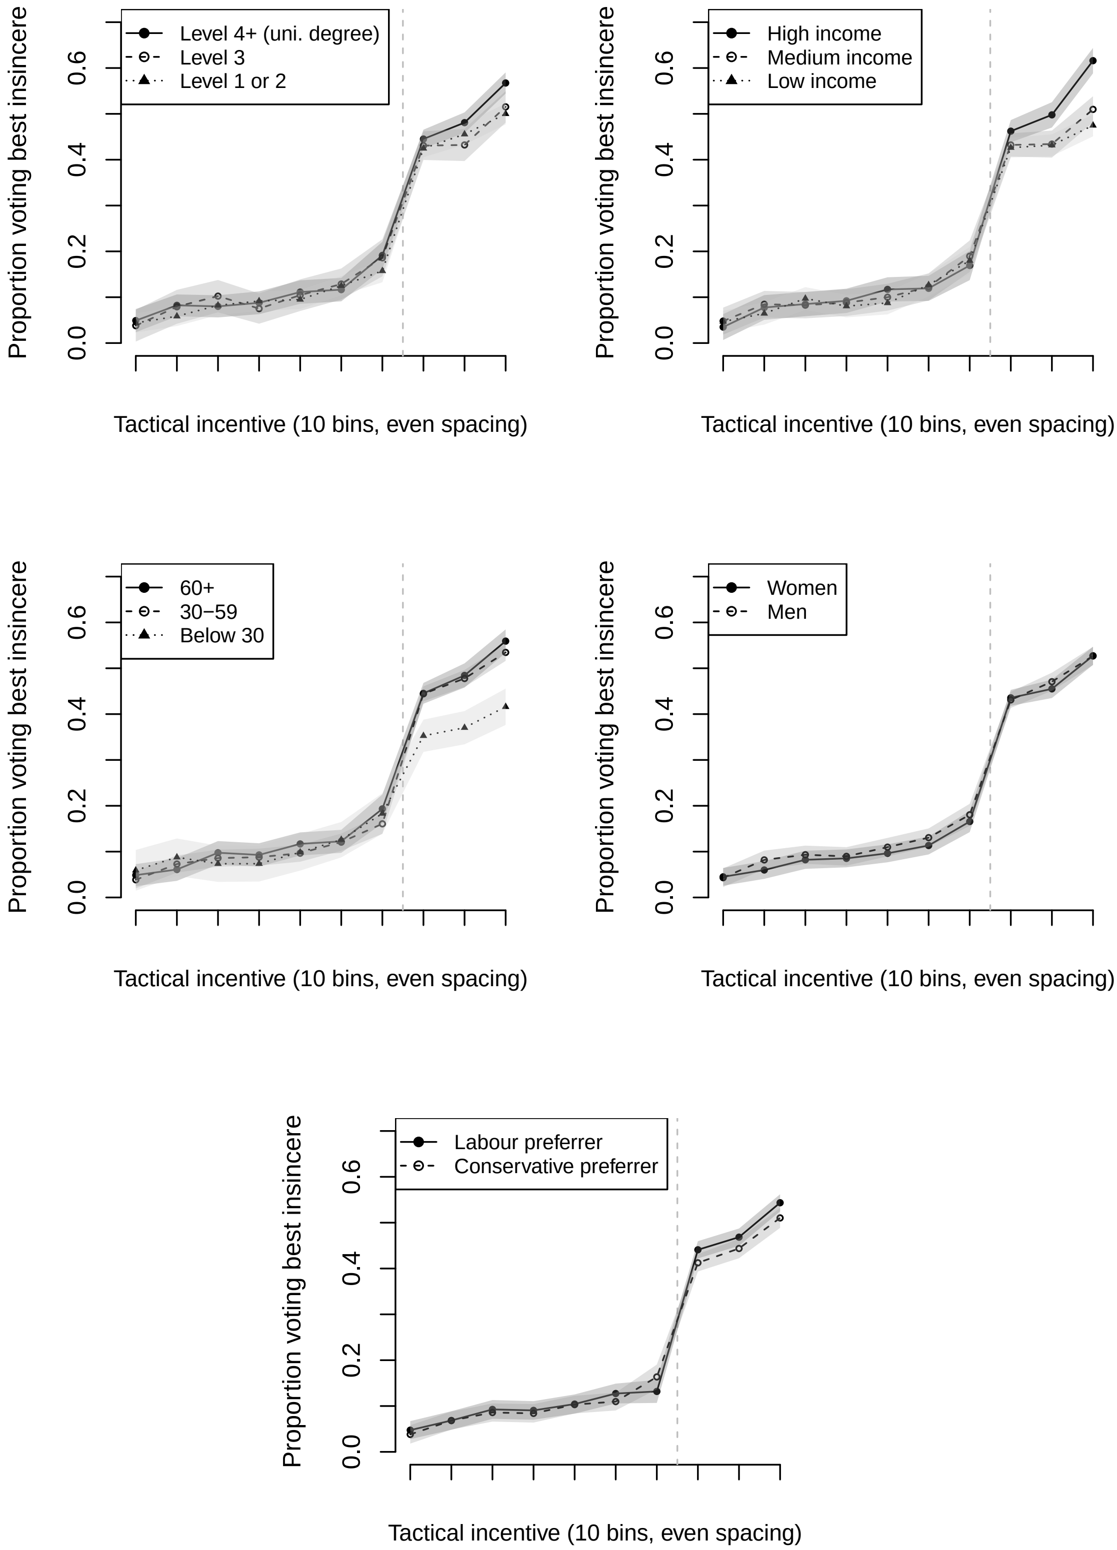

Note: Each diagram shows, as a function of τ and with subsetting by social characteristic, the proportion of respondents casting any insincere vote.

Figure C.9: **Alternative basis for the model of counterfactual election outcomes**: Model centered on forecasts rather than observed results

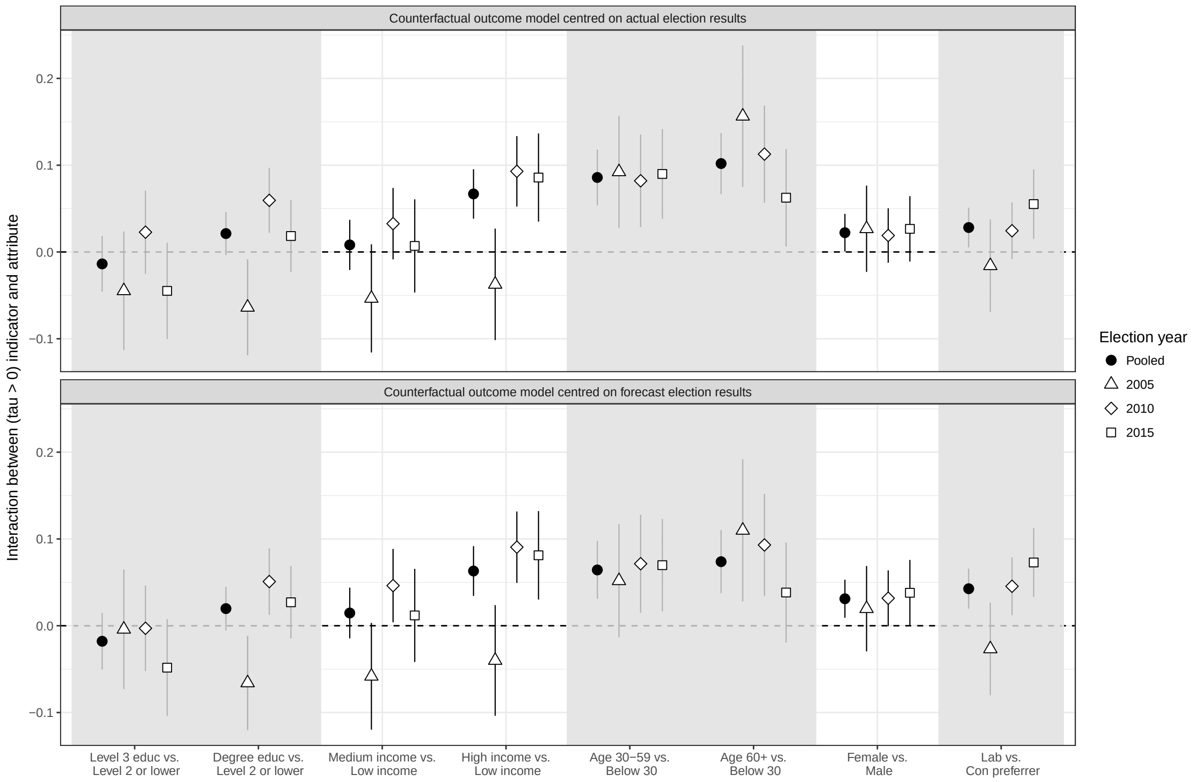

Figure C.10: **Greater uncertainty in model of counterfactual elections**: precision parameter set to $s = 12$ or $s = 20$ instead of $s = 85$

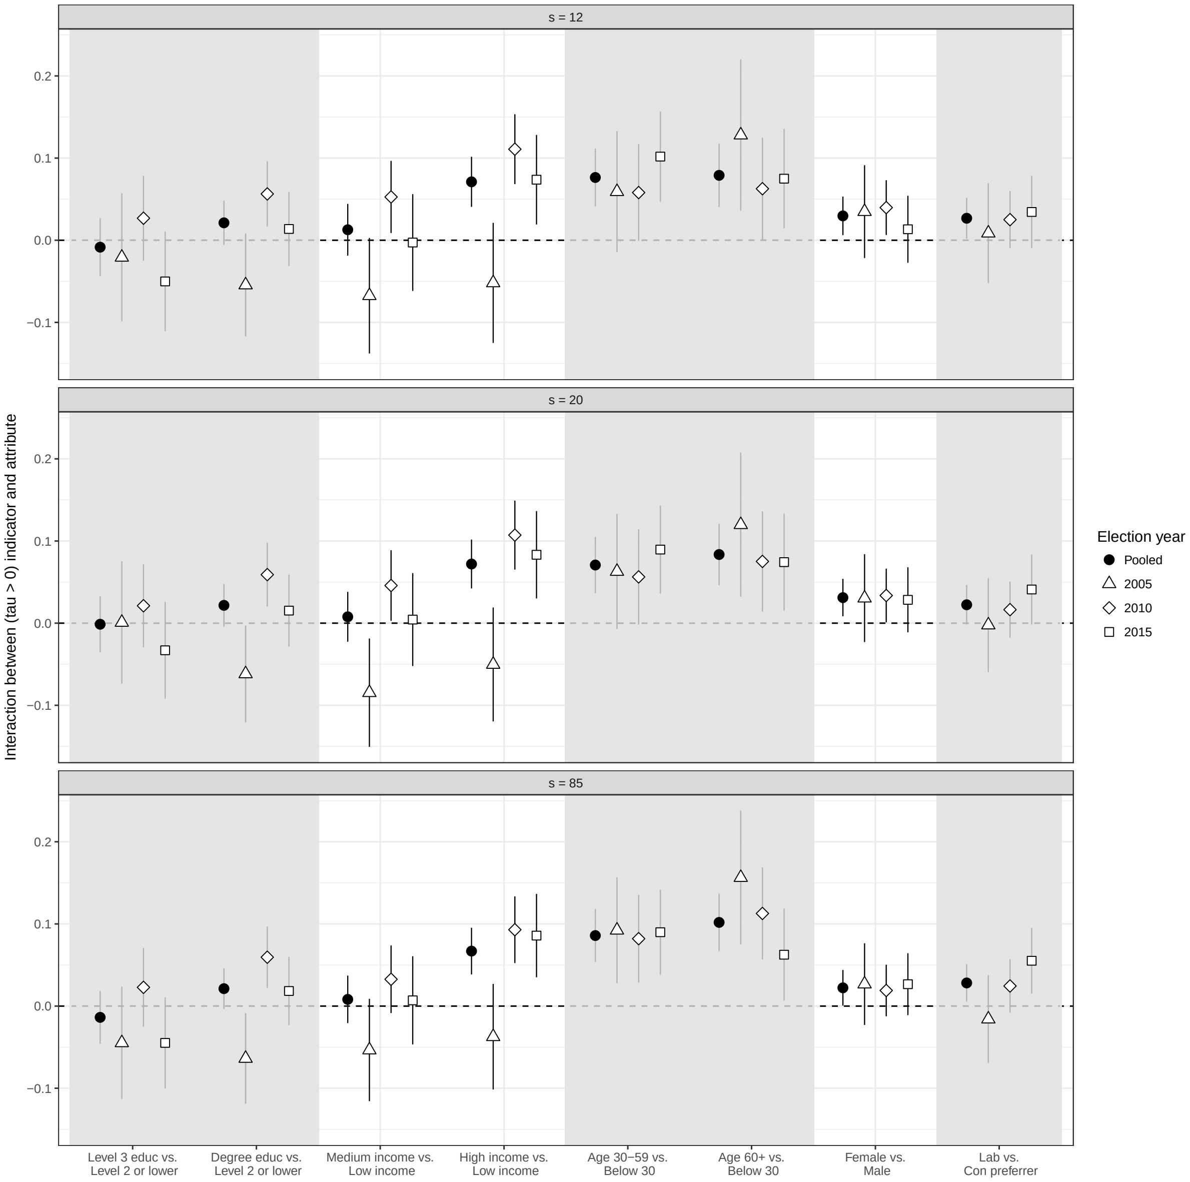

References

Evans, Geoffrey. 1994. Tactical Voting and Labour's Prospects. In *Labour's last chance? The 1992 election and beyond*, ed. Anthony Heath, Roger Jowell and John Curtice. Dartmouth pp. 65–84.

Fisher, Stephen D. 2001. "Extending the rational voter theory of tactical voting." Paper presented at the Annual Meeting of the Midwest Political Science Association.

Fisher, Stephen D. and David P. Myatt. 2017. "Strategic Voting in Plurality Rule Elections." Working paper.

Gelman, Andrew, Gary King and W John Boscardin. 1998. "Estimating the probability of events that have never occurred: when is your vote decisive?" *Journal of the American Statistical Association* 93(441):1–9.

Gelman, Andrew, Nate Silver and Aaron Edlin. 2012. "What is the probability your vote will make a difference?" *Economic Inquiry* 50(2):321–326.

Herrmann, M, S. Munzert and P. Selb. 2015. "Determining the effect of strategic voting on election results." *Journal of the Royal Statistical Society Series A* .

Hoffman, Dale T. 1982. "A model for strategic voting." *SIAM Journal on Applied Mathematics* 42(4):751–761.

Myatt, David P. 2000. "Multi-party tactical voting incentives." *Nuffield College Economics Discussion Paper* p. W19.

Niemi, R.G., G. Whitten and M.N. Franklin. 1992. "Constituency characteristics, individual characteristics and tactical voting in the 1987 British General Election." *British Journal of Political Science* 22:229–240.

Palfrey, Thomas R. 1989. *Models of Strategic Choice in Politics.* Ann Arbor: Michigan University Press chapter A mathematical proof of Duverger's law.