

STRATEGIC VOTING EXPERIMENTS

STEPHEN D. FISHER AND DAVID P. MYATT

NUFFIELD COLLEGE AND ST. CATHERINE'S COLLEGE, UNIVERSITY OF OXFORD

`stephen.fisher@nuf.ox.ac.uk` and `david.myatt@economics.ox.ac.uk`

Nuffield College Politics Working Paper 2002-W4

ABSTRACT. We report preliminary results from voting experiments designed to test both decision-theoretic and game-theoretic models of strategic voting. The experiments used a qualified majority voting system in which either of two options required a two-thirds majority to win, and thus for all participants to receive a positive payoff. Since the payoffs varied across the electorate, there were opportunities for strategic voting. An individual might choose to vote for the option that would give her a lower payoff, to ensure that the qualified majority was achieved. Voters were provided with samples of payoffs (or opinion polls) to provide a signal of the support for each of the two options. Some of these were provided to everyone publicly, others to individuals privately. The results so far are consistent with the comparative statics of a formal decision-theoretic model of strategic voting which predicts that the public and private information are of equal importance. We reject a full game-theoretic model since voters did not attach greater weight to the public information.

1. OVERVIEW

The phenomenon of strategic voting has long been of interest to political scientists. Under a variety of electoral systems, individuals may choose to vote for someone other than their preferred candidate. This is particularly relevant in the context of the “first past the post” plurality rule. An informal and familiar argument is as follows. A supporter of a trailing candidate may believe that this candidate has little chance of competing for the lead in the election. Any vote for such a candidate may well be “wasted.” Such an individual may then switch her vote to one of the leading candidates in the hope of exerting greater influence on the outcome of the election. This illustrates the central trade-off faced by any instrumental

We are very grateful to the British Academy for funding this research, and to Nuffield and St. Catherine's College for the use of their facilities. Also, we would like to thank David Austen-Smith, Jim Engle-Warnick, Shepley Orr, David Sanders, Hyun Shin, Charlotte Vinnicombe and Peyton Young for their help, comments and support. This paper was previously presented to the Annual Meeting of the American Political Science Association, San Francisco, August-September 2001.

voter. She must balance her relative preference for the different candidates against the relative likelihood of influencing the outcome of the election.

The determinants of strategic voting are of direct interest to political scientists. Furthermore, researchers have also sought to investigate the effect of strategic voting on political party systems. Strategic voting is an important component of Duverger's (1954) Law, which claims that "the simple-majority single-ballot system favours the two-party system." Subsequent authors have strengthened this claim. Palfrey (1989) claims that "... with instrumentally rational voters and fulfilled expectations, multicandidate contests under the plurality rule should result in only two candidates getting any votes." In other words, the pressure to vote strategically is so strong that all instrumental voters will vote for one of the two leading candidates.

In related work, however, it has been shown that the results of these formal models relied upon a statistical independence assumption (Myatt 2000a). Authors of many formal models assume that voting decisions are drawn independently from a commonly known distribution. Whereas the realised voting decision of an individual is unknown, the *aggregate* voting behaviour, at least in a large electorate, is certain. This means that all voters are perfectly aware of the "state of play" in a constituency. It is this feature that generates the stark Duvergerian outcomes with only two candidates receiving votes.

Why is this? We must recall that a single vote may only affect the outcome of an election if it is "pivotal"; in other words, there must be a tie for the lead. An instrumental voter will then compare the relative likelihood of different pivotal outcomes before casting her vote. Now, as the size of an electorate grows large, the absolute probability of a pivotal outcome falls away to zero. Importantly, however, it is the *relative* likelihood that matters. In a plurality rule election, the probability of a pivotal outcome involving the leading two candidates becomes infinitely greater than a tie involving any other pair. We can express this property in another way. In a large electorate, *conditional* on a pivotal outcome occurring, the tie will *almost always* be between the two leading candidates. It follows that any instrumental voter will cast their vote for one of this pair.

This does not happen, however, when the statistical independence assumption is relaxed. Suppose, for instance, that an individual is uncertain of the state of play in a constituency. We can formalize this by supposing that the voter is uncertain of the probability with which a randomly selected individual will vote for a particular candidate. This seems very reasonable — it would be rare indeed for a voter to be certain of the constituency-wide support for a particular candidate. Indeed, such certainty would allow the voter to predict the outcome of the election perfectly! More subtly, however, this uncertainty ensures that votes are no longer statistically independent. The voting decision of one individual reveals information about the

support levels for the different candidates, and hence the likely voting decision of further individuals. Furthermore, in a large electorate, this ensures that the incentive to vote strategically remains finite. A finite strategic incentive may change in response to electoral parameters. This permits comparative statics exercises, generating potentially testable predictions.

In related work (Myatt 2000a, Fisher 2000) we generate and test such comparative statics in the context of plurality voting. Myatt (2000a) builds a decision-theoretic model of a three-candidate plurality election. Voters obtain information by sampling the voting intentions of others in the electorate. In common with the informal intuition of Cain (1978) and others, the incentive to vote strategically increases with the distance from contention of a voter's preferred candidate. In contrast to the informal intuition, however, the incentive to vote strategically also rises with the margin of victory when the distance from contention is fixed. These comparative statics are confirmed by the empirical work of Fisher (2000). A second comparative static is also available from the theoretical model. When a preferred candidate trails in third place, the incentive to vote strategically should increase with the precision of information available to the voter. Unfortunately, we are unable to test this prediction based on survey data. It is the desire to assess this prediction that leads us to consider a series of experimental tests.

Our discussion so far has focused on a decision-theoretic setting. We consider strategic voting to be decision-theoretic when individuals act optimally assuming that others in the electorate will vote sincerely, or at least in accordance with their stated voting intentions. Of course, an intelligent voter might be expected to anticipate strategic voting behaviour by others. To cope with this, we require a game-theoretic perspective, as taken by authors such as Cox (1994) and Myerson and Weber (1993). This leads to the possibility of self-reinforcing strategic voting. An initial perceived bias in favour of one particular candidate may lead to some strategic voting away from less favoured candidates. But this strategic switching increases the distance from contention of the latter candidate, and hence increases the incentive for others to vote strategically. This process might continue until all support for a trailing candidate is lost in equilibrium — a strictly Duvergerian outcome.

The self-reinforcing logic described above is misleading. Notice that we referred to “the trailing candidate.” An implicit assumption, therefore, is that the identity of a trailing candidate is known. This is not the case when voters base their opinion of the strength of candidate support on privately observed information sources. One voter may receive a strong signal in favour of one candidate, whereas another will receive a strong signal in favour of a different candidate. This issue is addressed in Myatt (2000b). Perhaps surprisingly, when voters are privately informed, strategic voting exhibits negative rather than positive feedback. In other words, an increase in the tendency to vote strategically by others actually *reduces* the incentive for an individual

to vote strategically. Informally, an individual will become increasingly concerned that they may be switching their vote in the wrong direction. This does not happen when voters base their decisions on public sources of information. Each voter can be assured that all others are observing the same information, and hence can coordinate on an appropriate candidate. Summarizing, it is not just the precision of information sources that matter, but their nature. The theoretical prediction of Myatt (2000b) is that a publicly observed signal will have a much greater effect than a privately observed signal. We wish to test this prediction. To do so, however, we must consider an experimental setting.

Rather than consider a direct model of plurality voting, we follow earlier work and consider a simplified model of qualified majority voting. This captures all the salient aspects of strategic voting, but in a simpler framework. Furthermore, it allows for a more direct theoretical foundation for the experimental design. The experiment itself required each participant to cast a vote for one of two options A and B. A two-thirds majority was required for either of them to win and for all participants to receive a positive payoff. If neither received two-thirds of the vote, then no payoffs were received. The experiment was repeated 18 times (6 times each for three groups of participants). Each time the participants were presented with a random sample of payoffs (or opinion poll) which they could use to gauge the preferences of the other participants. These were sometimes private and individual, sometimes public and visible on an overhead projector, and sometimes there was both a private and a public sample. Some voters chose to vote for the option yielding the lower payoff. We interpret this as strategic voting. The pattern of strategic voting is consistent with the decision-theoretic model presented here. We reject the game-theoretic model since there is no evidence that voters placed more weight on the public information than the private information.

The structure of the remainder of this working paper is as follows. We describe the model in Section 2, and illustrate our ideas using the 1970 New York senatorial election. We review the optimal behaviour of a strategic voter in Section 3, and generate more specific predictions for behaviour via more careful modelling of voter preferences and information sources in Section 4. This leads directly to our experimental design in Section 5. The results of the experiments are reported in Section 6. Finally, we conclude in Section 7.

2. A QUALIFIED MAJORITY VOTING GAME

In this section, we review a simple framework that is able to generate strategic voting. The key driving force behind strategic voting is the tension between a individual's own personal preferences and the effectiveness of her vote. Indeed, her vote may well be most effective when it is "coordinated" with fellow members of the electorate. We motivate this idea by reviewing

TABLE 1. 1970 New York Senatorial Election

Candidate	Votes	Share
James Buckley	2,288,190	39%
Charles Goodell	1,434,472	24%
Richard Ottinger	2,171,232	37%
Total	5,893,894	100%

(a) Election Results

Parameter	Value	Comment
j	0	Buckley
j	1	Ottinger
j	2	Goodell
$n + 1$	3,605,704	Liberal voters
$\gamma_n n$	2,288,190	Disliked conservative voters
γ	0.634603	Required majority of liberal voters

(b) Stylized Parameterization

the familiar example of the 1970 New York senatorial election, and then formalize this as a stylized qualified majority voting game.

2.1. The 1970 New York Senatorial Election. The 1970 New York senatorial election was highlighted by Riker (1982) and others as an example of a “non-Duvergerian” outcome of strategic voting. In a “three horse race” two liberal candidates, Richard L. Ottinger and Charles E. Goodell, competed against the conservative James R. Buckley. More specifically, Goodell was an incumbent Republican who had taken a liberal stance on the Vietnam War, and hence received the nomination of the Liberal Party. The New York Conservative Party, however, rather than nominating Goodell as a “fusion” candidate instead supported the conservative Buckley.¹ We present the outcome of this election in Table 2(a). A widely held belief was that the liberal vote was split between Goodell and Ottinger, allowing the win for Buckley. Interestingly, we can view this situation as a qualified majority voting game among the liberal electorate. We can imagine that liberals, representing 61% of the electorate, all ranked Buckley as their least preferred candidate. In order to defeat the disliked conservative, 39% of the electorate must vote for one of the liberal candidates. This requires $39\%/61\% = 64\%$ of the liberal voters to coordinate on either Goodell or Ottinger. In other words, it is *not* enough for the liberal voters to achieve a simple majority among themselves for one liberal candidate. Rather

¹This is a traditional example of a three horse race and is used effectively in the undergraduate text of Morton (2001).

a *qualified* majority (in this case 64%) of liberals is needed to avoid the disliked Conservative nominee.

2.2. Qualified Majority Voting. We may formalize this idea by reviewing a system of qualified majority voting, as described by Myatt (2000b), from which this specification is taken. There are $n + 1$ voters, indexed by $i \in \{0, 1, \dots, n\}$. A collective decision is taken via qualified majority voting. Specifically, there are three possible actions $j \in \{0, 1, 2\}$, where $j = 0$ represents the (disliked) status quo. Each individual must cast a single vote for either of the two options $j \in \{1, 2\}$. Denoting the vote totals for each of these options by x_1 and x_2 respectively, it follows that $x_1 + x_2 = n + 1$. Based on these votes, the action implemented is:

$$j = \begin{cases} 0 & \max\{x_1, x_2\} \leq \gamma_n n \\ 1 & x_1 > \gamma_n n \\ 2 & x_2 > \gamma_n n \end{cases} \quad \text{where} \quad \gamma_n = \frac{\lceil \gamma n \rceil}{n} \quad \text{and} \quad \frac{1}{2} < \gamma < 1$$

The restriction $\gamma > 1/2$ ensures that first, it is impossible for both options 1 and 2 to meet the winning criterion of $x_j > \gamma_n n$, and second, the winning option must have a strict majority of the $n + 1$ strong electorate in order to win. The parameter γ gives a measure of the degree of coordination required to implement one of the actions $j \in \{1, 2\}$. For $\gamma \downarrow \frac{1}{2}$, only a simple majority is required. For $\gamma \uparrow 1$, complete coordination of the electorate is needed to avoid the status quo.

Payoffs are contingent only on the implemented action. The payoff u_{ij} is received by individual i when action j is implemented. All voters strictly prefer both outcomes $j \in \{1, 2\}$ to the status quo. This yields the payoff normalization $u_{i0} = 0$ and hence $\min\{u_{i1}, u_{i2}\} > 0$. The relative preference for the two options varies throughout the electorate. As we shall see, the log ratio $\tilde{u}_i \equiv \log[u_{i1}/u_{i2}]$ is sufficient to determine an individual's preferences.

We can immediately relate this formal specification back to the New York senatorial election. A stylized interpretation would be a zero payoff for a win by Buckley and positive payoffs for wins by either of the other candidates. We give a specific parameterization in Table 2(b). Other stylized illustrations of this formal framework are also available. In the context of the 1997 United Kingdom General Election, candidates $j \in \{1, 2\}$ might correspond to the Labour and Liberal Democrat parties, whereas candidate $j = 0$ might correspond to the Conservative party.

3. VOTING BEHAVIOUR IN LARGE ELECTORATES

In this section, we wish to characterize optimal voting behaviour in a large electorate. The analysis is drawn directly from Myatt (2000b), and readers are referred to that paper for proofs of the appropriate propositions.

3.1. Optimal Voting Behaviour. We begin by considering the behaviour of individual $i = 0$. In the familiar way, she may only influence the outcome of the election when there is a pivotal situation. To formalize this, we write x for the total number of votes cast for option 1, excluding the vote of $i = 0$. If $x = \gamma_n n$, then an additional vote will implement option 1 rather than the status quo. Similarly, if $n - x = \gamma_n n \Leftrightarrow x = (1 - \gamma_n)n$, then a single vote will tip the balance to option 2. Conditioning on any information available to the focal voter $i = 0$, consider the behaviour of the remaining voters. We write:

$$q_1 = \Pr[x = \gamma_n n] \quad \text{and} \quad q_2 = \Pr[n - x = \gamma_n n]$$

Hence q_1 and q_2 are the *pivotal probabilities* for options 1 and 2, in which one more vote is required to implement each of these options. Voting for option 1 will turn the status quo into the implementation of action 1 with probability q_1 , and yield an expected payoff of $q_1 u_1$, relative to abstention. Similarly, a vote for option 2 has expected payoff $q_2 u_2$.² Although our formal specification rules out abstention, it is clear that some vote is optimal whenever $\min\{q_1, q_2\} > 0$. Summarizing, optimal voting behaviour must satisfy:

$$\text{Vote} \begin{cases} 1 & q_1 u_1 > q_2 u_2 \\ 2 & q_2 u_2 > q_1 u_1 \\ 1 \text{ or } 2 & q_1 u_1 = q_2 u_2 \end{cases}$$

Before proceeding with our analysis, we wish to focus on the case where pivotal outcomes for both options are possible, so that $\min\{q_1, q_2\} > 0$. We are then able to define the *pivotal log likelihood ratio* or *strategic incentive* as $\lambda = \log[q_1/q_2]$. Employing this definition, the optimal decision rule when $q_1 u_1 \neq q_2 u_2$ becomes:

$$\text{Vote 1} \Leftrightarrow q_1 u_1 > q_2 u_2 \Leftrightarrow \log \left[\frac{u_1}{u_2} \right] + \lambda > 0 \Leftrightarrow \tilde{u} + \lambda > 0$$

The first term \tilde{u} represents the relative preference of a voter for option 1 versus option 2. Indeed, as anticipated in Section 2, this ratio is a sufficient description of a voter's preferences. The second element, and the key statistic of interest, is λ , the pivotal log likelihood ratio. This represents the relative influence of a vote for each of the two options, and is a convenient

²When considering the behaviour of the focal voter $i = 0$ we omit the subscript i for simplicity.

measure of the strategic incentive to switch to option 1. When $\lambda = 0$, there is no strategic incentive and voter $i = 0$ may support her preferred option.

3.2. Strategic Incentives and Voter Beliefs. We have seen that the log likelihood ratio of pivotal outcomes λ provides the incentive to vote strategically. We now ask what properties λ is likely to have. Our first observation is that a statistically independent model generates infinite strategic incentives in a large electorate. Suppose that each individual is expected to vote for option 1 with the independent probability p , and without loss of generality suppose that $p > 1/2$. The probability of a pivotal situation involving option 1 is then:

$$q_1 = \binom{n}{\gamma_n n} [p^{\gamma_n} (1-p)^{1-\gamma_n}]^n$$

with a similar expression holding for q_2 . Straightforward algebra then shows that:

$$\lambda = \log \left[\frac{q_1}{q_2} \right] = n \log \left[\frac{p^{\gamma_n} (1-p)^{1-\gamma_n}}{p^{1-\gamma_n} (1-p)^{\gamma_n}} \right] \rightarrow \infty \quad \text{as } n \rightarrow \infty$$

Of course, this property assumes that focal voter is certain of the parameter p . In other words, such an individual has clear and precise information on the “state of play” of the options in the qualified majority election.

Suppose instead that the voter is uncertain of p . We represent such uncertainty by the density $f(p)$, with full support $[0, 1]$. When calculating pivotal probabilities, a voter must take the uncertainty over p into account. Indeed, the pivotal probability q_1 becomes:

$$q_1 = \int_0^1 \binom{n}{\gamma_n n} [p^{\gamma_n} (1-p)^{1-\gamma_n}]^n f(p) dp \quad (1)$$

with a similar expression available for q_2 . Interestingly, as $n \rightarrow \infty$, the strategic incentive λ is completely determined by the density $f(p)$. The reason is as follows. As the constituency grows large, the Law of Large Numbers ensures that the realized fraction of the electorate voting for option 1 will be p . This means that the likelihood of a pivotal situation involving option 1 corresponds to the appropriate density evaluated at $p = \gamma$; in other words, $f(\gamma)$. Similar reasoning ensures that the log likelihood ratio of pivotal events will, in the limit, be equal to $f(\gamma)/f(1-\gamma)$. More formally, Myatt (2000b) shows that:

$$\lim_{n \rightarrow \infty} (n+1)q_1 = f(\gamma), \quad \lim_{n \rightarrow \infty} (n+1)q_2 = f(1-\gamma), \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{q_1}{q_2} = \frac{f(\gamma)}{f(1-\gamma)}$$

We wish to make a number of observations. First, the incentive to vote strategically remains finite, and hence comparative statics may be possible. Second, the strategic incentive is determined by uncertainty over support for the two options, represented by $f(p)$. We call this

“constituency uncertainty.” Third, we may build a microfoundation for this uncertainty, and hence generate appropriate predictions for use in both empirical and experimental settings.

4. MODELLING PREFERENCES AND SIGNALS

In this section, we describe an explicit model of voter preferences and information sources.

4.1. Voter Preferences. We begin by considering the log relative payoffs of a voter which we write as $\tilde{u}_i \equiv \log[u_{i1}/u_{i2}]$. We decompose this into *common* and *idiosyncratic* components as follows:

$$\tilde{u}_i = \log \left[\frac{u_{i1}}{u_{i2}} \right] = \eta + \epsilon_i$$

The component η is common to all individuals. In contrast, the *idiosyncratic* component ϵ_i is distributed independently across individuals, with distribution $\epsilon_i \sim N(0, \xi^2)$. In this sense, the parameter ξ^2 measures the idiosyncrasy of individuals in the electorate. The common component to each voter’s preferences corresponds to the support level for option 1. To see this, we may directly calculate the support level:

$$p = \Pr[\tilde{u}_i \geq 0] = \Phi \left(\frac{\eta}{\xi} \right)$$

In order to generate uncertainty over the support for the two options, we suppose that η is unknown to an individual voter. An individual may then consult public and private information sources to generate beliefs over η and hence p , yielding the appropriate strategic incentive.

4.2. Public and Private Signals. We suppose that individuals begin with a common and diffuse prior over η , with no knowledge of the common utility component. Information on η is then gleaned from two sources: *public* and *private* signals. We model these signals as opinion polls. These opinion polls are particularly detailed in that we suppose that they encompass intensity of preference for the candidates, as well as the identity of a preferred candidate. As explained in Myatt (2000b), we view a private opinion poll as societal communication between voters prior to an election.

To formalize this idea, we suppose that a voter i observes a sample $\{\tilde{u}_k\}$ of preferences drawn independently from the same distribution as her own. More concretely, we suppose that $\tilde{u}_k | \eta \sim N(\eta, \xi^2)$ iid. The total sample size is $M + m + 1$. The first component is a subsample of size M that is publicly observed. This means that all voters observe this subsample simultaneously. This corresponds to a public opinion poll of relative support for the two options. The second component is a subsample of size m that is privately observed. In other words, each voter observes an independently generated private sample of this size. This component corresponds to private societal communication — individual voters gain a signal of the support for the

options from their day to day lives. Finally, the last component is the voter's own preferences. She may regard her own relative preferences as a signal of η .

Before considering the behaviour of an instrumental voter, we may assemble the sampled preferences into appropriate sufficient statistics. Take, for instance, the public sample $\{\tilde{u}_k\}_{k=1}^M$, where $\tilde{u}_k | \eta \sim N(\eta, \xi^2)$ iid. Since each sample point is normally distributed, then the sample mean is sufficient for η . We can form the *public signal*:

$$\hat{\eta}_M = \frac{1}{M} \sum_{k=1}^M \tilde{u}_k$$

Similarly, we can also form the *private signal*:

$$\hat{\eta}_m = \frac{1}{m} \sum_{k=M+1}^{M+m} \tilde{u}_k$$

Finally, we can use a standard Bayesian updating procedure to find the posterior belief of voter i over η . Assuming that she begins with a diffuse prior, her posterior belief satisfies:

$$\eta \sim N(\bar{\eta}, \sigma^2) \quad \text{where} \quad \bar{\eta} = \frac{M\hat{\eta}_M + m\hat{\eta}_m + \tilde{u}_i}{M + m + 1} \quad \text{and} \quad \sigma^2 = \frac{\xi^2}{M + m + 1}$$

Thus posterior beliefs over η are normally distributed with a mean that is linear in the public signal, private signal, and a voter's own log relative payoff.

4.3. Decision Theoretic Strategic Voting. Consider the decision of voter i . Suppose that she is willing to vote strategically, but anticipates that all other individuals will vote sincerely. We call this *decision-theoretic* strategic voting. This is when a voter acts strategically but fails to take into account the strategic behaviour of others.

Sincere voting by the remaining electorate means that a randomly selected individual will vote for option 1 with probability $p = \Phi(\eta/\xi)$. Of course, this probability must be evaluated conditional on η . From the perspective of our sophisticated voter i , η is unknown and hence p is uncertain. We must consider the cumulative distribution function $F(p)$, conditional on any information available to voter i . Straightforwardly, we have:

$$F(p) = \Pr \left[\Phi \left(\frac{\eta}{\xi} \right) \leq p \right] = \Pr \left[\eta \leq \xi \Phi^{-1}(p) \right] = \Phi \left(\frac{\xi \Phi^{-1}(p) - \bar{\eta}}{\sigma} \right)$$

where this last equality follows from the fact that the voter's posterior belief over η is normal with mean $\bar{\eta}$ and variance σ^2 . Of course, we may differentiate this expression to obtain $f(p)$.

Evaluating at γ and $1 - \gamma$, taking ratios and logs yields:

$$\begin{aligned} \log \left[\frac{f(\gamma)}{f(1-\gamma)} \right] &= \frac{2\xi\Phi^{-1}(\gamma)\bar{\eta}}{\sigma^2} = \frac{2\Phi^{-1}(\gamma)(M\hat{\eta}_M + m\hat{\eta}_m + \tilde{u}_i)}{\xi} \\ &= \frac{2\Phi^{-1}(\gamma)}{\xi} \left\{ \tilde{u}_i + \sum_{k=1}^{M+m} \tilde{u}_k \right\} \end{aligned}$$

This term is, of course, an exact measure of the strategic incentive faced by the voter. The optimal voting rule becomes:

$$\begin{aligned} \text{Vote 1} &\Leftrightarrow \left[1 + \frac{2\Phi^{-1}(\gamma)}{\xi} \right] \tilde{u}_i + \frac{2\Phi^{-1}(\gamma)}{\xi} \sum_{k=1}^{M+m} \tilde{u}_k \geq 0 \\ \text{or} &\left[1 + \frac{2\Phi^{-1}(\gamma)}{\xi} \right] \tilde{u}_i + \frac{2\Phi^{-1}(\gamma)}{\xi} (M\hat{\eta}_M + m\hat{\eta}_m) \geq 0 \end{aligned}$$

This gives us a prediction for voter behaviour within the context of this model, assuming that individuals behave decision-theoretically. Notice that the voter responds equally to components of the public and private opinion polls. Furthermore, she responds more strongly to her own log relative preference than to other sampled preferences.

4.4. Game Theoretic Strategic Voting. The analysis above models voters as decision-theoretic; they do not account for the possibility of strategic behaviour by others. To extend the analysis further, we must allow voters to take into account such strategic behaviour. Such a game-theoretic perspective is taken by Myatt (2000b). He considers a class of symmetric and monotonic voting strategies. A voting strategy is *symmetric* if an individual's vote depends only on her preferences and private signal, and not on any payoff-irrelevant factors such as the particular identity of the voter. A strategy is *monotonic* if an increase in \tilde{u}_i or the private signal increase the probability of a vote for the first option.

Restricting to this class of strategies, Myatt (2000b) obtains the following results. If all other voters employ a symmetric strategy, then it is a best response for an individual voter i to employ a linear strategy of the following form:

$$\text{Vote 1} \Leftrightarrow \underbrace{\tilde{u}_i + \alpha\hat{\eta}_M + \beta \left[\frac{m\hat{\eta}_{mi}}{m+1} + \frac{\tilde{u}_i}{m+1} \right]}_{\lambda_i = \text{strategic incentive}} + \psi \geq 0$$

Here we have written $\hat{\eta}_{mi}$ to emphasize that this is the private signal received by voter i . By inspection, we see that the strategic incentive faced by a voter is a linear function of the public and private signal. An easy corollary is that the class of linear strategies is closed under best response. If all other voters use a symmetric monotonic linear strategy with parameters α , β and ψ , then it is a best response for an individual to use a linear strategy with parameters $\hat{\alpha}$,

$\hat{\beta}$ and $\hat{\psi}$ where:

$$\begin{aligned}\hat{\alpha} &= \frac{\hat{\beta}}{1+\beta} \frac{\alpha(M+m+1) + (1+\beta)M}{m+1} \\ \hat{\beta} &= \frac{2\Phi^{-1}(\gamma)}{\xi} \frac{\sqrt{(m+1)(m+1+\beta^2+2\beta)}}{1+\beta} \\ \hat{\psi} &= \frac{\hat{\beta}}{1+\beta} \frac{M+m+1}{m+1} \psi\end{aligned}$$

Notice that both $\hat{\alpha}$ and $\hat{\psi}$ are increasing in α and ψ respectively. This is the self-reinforcing aspect of strategic voting. In contrast, $\hat{\beta}$ is *decreasing* in β . As others in the electorate increase their responses to their private signals, it is optimal for an individual to *reduce* her response in turn. Myatt (2000b) explains this self-attenuating aspect of strategic voting in details.

To find an equilibrium, we may look for fixed points of the best response mappings given above. For $\hat{\beta}(\beta)$ this is straightforward. This has a unique positive and stable solution β^* , satisfying:

$$\beta^* = \frac{2\Phi^{-1}(\gamma)}{\xi} \frac{\sqrt{(m+1)(m+1+(\beta^*)^2+2\beta^*)}}{1+\beta^*} \quad (2)$$

Equilibrium solutions for α and ψ are more problematic. Notice by inspection that both have interior solutions. Trivially $\psi^* = 0$ solves the final equation, and the fixed point α^* satisfies:

$$\alpha^* = \frac{\beta^*(1+\beta^*)M}{m+1-\beta^*M}$$

There are alternative solutions, however. Setting $\alpha = \pm\infty$ is a possible solution — this corresponds to all individuals voting for a single option. Myatt (2000b) offers a way to select between these equilibria by requiring stability. Briefly, one may start with an appropriate monotonic and symmetric strategy profile, and compute a sequence of best responses. Such a sequence converges to the interior equilibrium if and only if:

$$m+1 > \beta^*M$$

Informally, this inequality says that private information is relatively important compared to public information. If this inequality is not satisfied, then an iterative best response process converges to an extreme equilibrium where all voters perfectly coordinate on a single option. In addition, the above inequality is required to ensure that $\alpha^* > 0$ is satisfied. This means that a voter responds positively rather than negatively to the public signal. Informally, this means that a public opinion poll in favour of an option helps strategic voting toward it rather than away from it.

4.5. Summary. In summary, our predictions are as follows. If voters are expected to formulate their votes in a decision-theoretic way, then we would expect:

$$\text{Vote 1} \Leftrightarrow \left[1 + \frac{2\Phi^{-1}(\gamma)}{\xi} \right] \tilde{u}_i + \frac{2\Phi^{-1}(\gamma)}{\xi} \sum_{k=1}^{M+m} \tilde{u}_k \geq 0 \quad (3)$$

In a game theoretic-world there are two possibilities. First, we calculate the coefficient β^* from Equation 2. Then we have:

$$m + 1 < \beta^* M \quad \Rightarrow \quad \text{Full Coordination}$$

However, if $m + 1 > \beta^* M$ then we would expect:

$$\text{Vote 1} \Leftrightarrow \left[1 + \frac{\beta^*}{m+1} \right] \tilde{u}_i + \frac{\beta^*(1+\beta^*)}{m+1-\beta^*M} \sum_{k=1}^M \tilde{u}_k + \frac{\beta^*}{m+1} \sum_{k=M+1}^{M+m} \tilde{u}_k \geq 0$$

In this latter case, the attention paid to the public signal is much greater than that paid to the private signal.

5. EXPERIMENTAL DESIGN

5.1. Implementing the Theory. The theoretical models of the companion papers and the previous section make a number of predictions. They calculate the incentive to vote strategically as a function of the information sources on which a voter might base her decision. Some of the comparative static predictions may be assessed using empirical data. In the plurality rule analog of the model above, changing the state of play of a constituency will change the signals received by voters and hence their behaviour (Myatt 2000a, Fisher 2000). There are some features, however, than simply cannot be addressed using existing survey data. Notice that we may vary the amount and type of information given to voters in the theoretical framework. Perhaps the only way to assess these aspects is to conduct a series of experiments.

Of course, we cannot implement the theoretical model directly. There are a number of reasons. First, the model is specified using von Neumann Morganstern utilities as payoffs. In an experimental setting, we are only able to generate monetary payoffs. Second, the model assumes that voters are aware of the data generating process. Whereas there is no knowledge of η , voters know that conditional on η we have $\tilde{u} \sim N(\eta, \xi^2)$. In particular, they are aware of the parameter ξ , which is a measure of the heterogeneity of payoffs. Third, the theoretical predictions hold only when the electorate is sufficiently large. They hold only approximately for finite electorates.

Nevertheless, we wish to test some of the qualitative predictions of the decision-theoretic and game-theoretic models. Roughly speaking, the three hypotheses that we may consider as follows.

- (1) In a decision-theoretic world, public and private samples of voter preferences should have a similar effect on strategic voting.
- (2) In a game-theoretic world, when public information is sufficiently precise (or equivalently $m + 1 < \beta^* M$), then there should be full or nearly full coordination.
- (3) In a game-theoretic world in which public information is relatively less precise, the public signal should have a greater effect than the private signal.

It is the first two of these hypotheses that we have tested in the experiments so far.

5.2. Recruitment and Organisation. We ran three experimental sessions in late May 2001 in Oxford. Recruitment was mainly by emails sent to various people in the Oxford area, but also by posters displayed in colleges. Inevitably most of the participants were students, both graduate and undergraduate from across the university, but there were also several young to middle-aged professionals. The invitation made it clear that they would receive at least £5 and possibly up to £20 in cash if they completed the hour long session of mock experiments (\$7 and \$28 respectively). For most people these are reasonable amounts worth participating for. We also emphasized that this had nothing to do with the general election and no knowledge of politics was required.

On arrival, participants registered their name and address and were given an instruction sheet and asked to sit at a desk suitably spaced apart from other participants. Once everyone had arrived the instructions were read to the participants and questions of clarification were answered. The instruction sheet included practice questions which the participants all answered and these were checked by us for accuracy. The questions were designed to make the participants think carefully about the structure of the elections and wrong answers were usually due to haste rather than misunderstanding. Answers were explained to those who got them wrong and the elections did not start until all the participants were satisfied that they understood how they would work. The instruction sheet, including the procedure for the elections, was as follows.

5.3. Instruction Sheet for Participants. Thank you very much for coming here today and participating in these voting experiments. We appreciate the time and effort you have made to help us with our research.

You will receive a £5 participation fee if you complete the session and follow the instructions. Most importantly you must not talk or communicate in any way with other participants. Also

you must not look at any other ballot papers except your own. If you break these rules we will ask you to leave and you will not be paid.

You will be asked to cast a vote in each of six elections today. Each of the elections will take the same form. There will be two options, A and B. Option A wins the election if two thirds or more of the votes are for option A. Likewise, option B wins if two thirds or more of the votes are for option B. If neither A nor B receives two thirds or more then neither wins.

For each election you will be given a ballot paper drawn at random. The ballot paper tells you, how much money you will receive if option A wins, and how much money you will receive if option B wins. We call these amounts ‘payoffs’. The payoffs are not the same for all participants. Rather they have been drawn at random from a large population. As an indication of what that population looks like, you will be given a random sample of the payoffs.

You must vote for either A or B by marking the relevant box. You must also write your participant number in the space available on the ballot paper. After voting in all six elections is over, we shall select one of the elections at random and count the votes to determine which option has won if either. If there is a winner, you will be paid the amount specified on your ballot paper for that election, in addition to the £5 participation fee. If there is no winner you will only receive the participation fee.

Finally, we wish to point out that all the information we give you is true. There is no attempt to trick or mislead you. Please now answer the practice example questions, so that we can check that you understand how the experiment works. If you have any questions you should raise your hand at this stage.

Thank you very much again for your help.

Yours

Stephen Fisher (Nuffield College) and David Myatt (St. Catherine’s College)

Practice example

If A wins you get £4. If B wins you get £11.

Assuming that this is the election that is randomly selected, if you voted for B and A wins, how much will you get including your participation fee?

If B receives 60 per cent of the vote, how much will you get including your participation fee?

5.4. Parameter Assignment. At the start of each election, ballot papers were handed out to the participants at random and left face down on their desk. At this point we explained just how the random sample(s) of payoffs would be displayed. Either there was a public signal,

or a private signal or both. A public signal was a random sample of payoffs projected on an overhead projector so that everyone could see, and everyone could see that everyone else could see the sample. A private signal was a unique random sample of payoffs shown on the ballot paper which only the participant with that ballot paper could see. But everyone had the same size sample we told the participants this. In each case the sample was simply a list of pairs of payoffs for A and B generated in the same way as the participant payoffs. As described in Section 4.1, the log ratio of the payoffs for A and B were such that,

$$\log \left[\frac{u_{iA}}{u_{iB}} \right] \sim N(\eta, \xi^2).$$

Actual payoffs were fixed by applying the constraint $u_{iA} + u_{iB} = \pounds 15$ and rounding the amounts to the nearest ten pence.

Table 2 shows the choice of η , ξ , the size of the public sample and the size of the private sample for each round in each experimental session. Since the elections required a two thirds majority, $\gamma = 2/3$ throughout. The parameters were not chosen at random but according various criteria. We were not particularly concerned with the values of η and ξ . A pilot study suggested that the initial choice of 0 and 1 produced reasonable sets of payoffs that were neither too diffuse nor too one-sided. Variation in the private and public sample sizes was considered most important. Also, we were keen to present different ‘treatments’ (public signal only, private signal only and both) in roughly equal measure. However, the order of these treatments varied from session to session. Treatment order could influence behaviour and we wished to minimize the possibility of a systematic effect.

At each election participants were given as much time as they wanted to decide how to vote, but this was usually no more than a couple of minutes. In each session the election that was chosen at random was one that yielded a two-thirds majority, so the participants went home with between $\pounds 5$ and $\pounds 20$. Unsurprisingly, they were all keen to be contacted about any further experiments!

6. EXPERIMENTAL RESULTS

6.1. Overview of the Data. The experiments produced very high quality data. All the participants voted in all the elections in their session. There were 103 participants each voting in 6 elections yielding 618 votes in total. There were no spoilt ballots and only one ballot paper was missing the participant number. Almost exactly one third of the votes were for the option yielding the lower payoff. Such votes could be strategic, but they could also be the result of misunderstanding or a deliberate attempt to disrupt the experiment. However, there was only two instances where a participant voted against the ‘favourite’ (option with the highest

TABLE 2. Parameter Assignment

Session	Round	η	ξ	Private sample size	Public sample size
1	1	0.1	1.0	0	5
1	2	0.0	1.0	5	0
1	3	0.0	1.0	5	5
1	4	-0.1	1.0	0	4
1	5	0.0	1.0	0	7
1	6	0.1	1.0	5	10
2	1	0.0	1.0	6	0
2	2	0.1	0.9	0	7
2	3	0.0	1.1	5	8
2	4	-0.1	1.0	0	3
2	5	0.0	0.9	7	0
2	6	0.1	1.0	9	5
3	1	0.1	1.0	7	7
3	2	0.0	1.0	0	9
3	3	-0.1	0.9	8	0
3	4	-0.1	1.1	4	0
3	5	0.0	1.0	0	6
3	6	0.0	1.0	5	3

payoff) when all the information available suggested that the favourite was the most popular option. There is no evidence that any participant was persistently voting against their own interests. There is also no evidence that anyone was too keen in their willingness to vote for the most popular option. No one voted strategically all six times. A handful of participants voted strategically four times and in each case inspection of the payoffs and information involved suggests that the behaviour was reasonable. Likewise there is no reason to suppose that some participants were resolutely failing to vote strategically despite incentives to do so. In those cases where people never voted strategically the choices seem defensible in light of the payoffs and information.

6.2. Analysis of Strategic Voting. In addition to the more direct test of the model presented in the following subsection, there are a number of hypotheses that relate to the probability of strategically deserting the preferred option. Some of these are intuitive, some are informal interpretations of the model predictions and some concern the nature of the experimental setting. Intuitively we expect those for whom the payoff for one option is much larger than the other to be less likely to vote strategically for the lower payoff.³ Also strategic voting should

³Note that we cannot test this against the alternative hypothesis that it is the absolute value of the lowest (or highest) payoff available that matters to someone not the difference between them. Since the payoffs for any individual sum to £15, the payoff for the favourite is a direct linear function of the difference in payoffs.

TABLE 3. Probit Model of Strategic Voting

	Coeff	Std. Err.	p-value	95% CI
Difference in own payoffs for-against	-0.092	0.021	0.000	(-0.133, -0.052)
Public number for-against	-0.224	0.059	0.000	(-0.340, -0.108)
Private number for-against	-0.194	0.068	0.004	(-0.326, -0.061)
Total of public payoffs for-against	-0.027	0.008	0.001	(-0.042, -0.012)
Total of private payoffs for-against	-0.043	0.011	0.000	(-0.064, -0.023)
Constant	-0.315	0.118	0.008	(-0.546, -0.083)

be more likely as the evidence from the public and private signals suggest that the favourite is not the most popular option. Both these hypotheses accord with the predictions of the model, but the latter is more subtle. One can imagine at least two responses to the random samples of payoffs. First, one can look at the number of people in the sample preferring each option. Strategic voting should increase with the number in the sample for the favourite minus the number against. Second, one could add up all the payoffs for A and all those for B and look at the difference between them. This is equivalent to looking at the average of the differences in the payoffs, but weighting the sample by the number of people in it. Using this technique, strategic voting should increase as the difference in the total payoffs favours the one's own preferred option. These two methods of viewing the sample do not always suggest similar conclusions. Sometimes they do not even point in the same direction.

These informal comparative statics are represented by variables in the probit model of strategic voting shown in Table 3. Those concerning the number in the sample for/against the favourite and the total payoffs in the sample for/against the favourite, are computed separately for public samples and private samples. The unique insight of the Myatt (2000a) model presented here is that, in a game-theoretic world, public information should be weighted more strongly than private information. In a decision-theoretic context they should be equally weighted. The coefficients in Table 3 show that the informal comparative statics do all seem to work well, but there is no evidence to reject the hypothesis that the public and private information is weighted equally. This is clear from inspection of the confidence intervals in the final column.

Since the experiments are unusual and take some thought, it would not be surprising to find that people's behaviour changed over the course of the session. For instance as people become more familiar with the election procedure they may be more willing to vote strategically. We found no evidence for such a general learning effect. However we did find that participants were less likely to vote strategically in the first election of the session than later rounds, after controlling for the terms in Table 3. We also found that the participants in the first experimental session were more likely to vote strategically, but this may be because they experienced more public samples than the other two sessions. Such issues affecting the data need to be explored in

greater depth. However, it is worth noting that including terms for session or round effect do not change the coefficients of the substantively interesting variables in Table 3. The conclusion from this informal analysis remains the same. Participants seem to behave as if they were in a decision-theoretic, rather than a game-theoretic context.

6.3. A Formal Test of the Model. Both the game-theoretic and the decision-theoretic models can be rejected by an extremely strict test. The models produce exact predictions as to how each participant should have voted in each election. In those elections with a public sample, the incentive to vote strategically under the game-theoretic model was infinite. All the participants should have voted for the same option. This was not the case in any of the elections, with or without a public signal. On a strict interpretation this is enough to reject the game-theoretic version of the model.

TABLE 4. Classification Table for the Decision-Theoretic Model

<i>Vote</i>	Favourite	Strategic	Total
<i>Predicted</i>			
Favourite	56.80	10.19	66.99
Strategic	9.87	23.14	33.01
Total	66.67	33.33	100.00

A strict test of the decision-theoretic version requires inspection of the predictions of equation 3 on page 13. These have been tabulated against the actual vote, and Table 4 shows that the voting behaviour does not correspond to the model in 20 per cent of cases. Perhaps more problematic is that 30 per cent of those predicted to vote strategically actually voted sincerely. Again on a strict interpretation we must also reject the decision-theoretic model.

However, it is unreasonable to expect a model to predict behaviour perfectly. As mentioned in subsection 5.1, there are various reasons why we cannot replicate the model even in an experimental framework. Participants should have von Neumann-Morgenstern utilities as payoffs, not cash incentives. Participants should know the distribution of the utilities and their variance, conditional on the mean. Also, electorate should be large, but the three sessions had 38, 39 and 26 participants. Given that we expect some deviation from the model, it is possible to operationalise equation 3 in a standard probit setting to examine the important comparative statics. To do this, we define the following system of explanatory variables:

$$\begin{aligned} \text{utilityratio} &= \left(1 + \frac{2\Phi^{-1}(\gamma)}{\xi}\right) \log \left[\frac{u_{Ai}}{u_{Bi}}\right] \\ \text{publicsignal} &= \frac{2\Phi^{-1}(\gamma)}{\xi} \sum_{k=1}^M \log \left[\frac{u_{Ak}}{u_{Bk}}\right] \\ \text{privatesignal} &= \frac{2\Phi^{-1}(\gamma)}{\xi} \sum_{j=1}^m \log \left[\frac{u_{Aj}}{u_{Bj}}\right] \end{aligned}$$

Given their definitions, if the probit analysis fits the decision theoretic model then the coefficients of `utilityratio`, `publicsignal` and `privatesignal` should all be equal to each other, and the constant should be zero. The estimated probit model is shown in Table 5. All three variables are positive and statistically significant, which adds support to both the decision-theoretic and game-theoretic models. Broadly speaking, the behaviour of participants was sensitive to the individual payoffs and the information from the samples as predicted by the strategic voting models. Having reached this conclusion, the question is whether people were behaving game-theoretically or decision-theoretically? Inspection of the coefficients shows that they are of similar magnitude, as would be predicted by the decision-theoretic model. The overlapping confidence intervals show that there is insufficient evidence to reject the comparative statics of the decision-theoretic framework. This accords with the informal analysis shown above. Certainly, if there is a departure from the decision-theoretic model it is not in the direction of the game-theoretic model. If people were reasoning game-theoretically one would expect the coefficient of the public signal to be greater than that of the private signal. The opposite is the case.

TABLE 5. Probit Test of the Theoretical Hypotheses

	Coef.	Std. Err.	p-value	95% CI
<code>utilityratio</code>	0.432	0.043	0.000	(0.348, 0.517)
<code>publicsignal</code>	0.321	0.031	0.000	(0.261, 0.381)
<code>privatesignal</code>	0.442	0.050	0.000	(0.344, 0.540)
Constant	-0.033	0.064	0.599	(-0.158, 0.091)

7. CONCLUDING REMARKS

The analysis thus far is preliminary. In subsection 5.1 we mentioned reasons why the experimental design did not match the model specifications. In addition to potential imperfections in the organisation of the experiments, there may be framing effects, learning and various other

phenomena well known to experimental economists (Kagel and Roth 1995). Although we tested for a simple learning effect, it may still be the case that participants strategy evolved over the course of the experiment. In particular, people are likely to be affected by their own experience in previous rounds which may make them more or less likely vote strategically.

Whilst some adjustment for error may be possible in the analysis of the data presented here, further experimentation is also necessary for three main reasons. First, the game-theoretic model has only been tested under extreme conditions where full coordination was predicted. Second, the decision-theoretic model has not been tested under conditions where participants are constrained to respond in a decision-theoretic manner, i.e. where game-theoretic reasoning is impossible. Thirdly, the model predictions depended on asymptotic theory. Future experiments will correct for these features the research presented here.

Nonetheless, the comparative statics of the decision-theoretic model do seem to fit the data well. To this extent there is a relatively strong substantive conclusion. It is surprising that people did not respond more strongly to the public signals than the private signals. The overhead projector should have provided a focal point for the participants. In the case of conflicting public and private signals, one would expect the public signal to be dominant. However, this did not happen. Rather than considering what other people might be experiencing and thinking, it seems that people behaved as if everyone else was going to vote for their favourite option.

This conclusion in favour of the decision-theoretic over the game-theoretic model has important implications for understanding elections. If an election is dominated by public signals, then there is considerably less strategic coordination than we would have expected under the game-theoretic model. The opposite is true when public information is scarce. In the case of a US presidential or senatorial contest there are well publicised opinion polls that could be said to constitute public information. However, in a British general election, little is known publicly about the state of play in any particular constituency since there are very rarely opinion polls published at that level. Instead people rely on the previous result (if they know it) and various private sources of information. The experimental results here suggest that strategic voting is lower in high profile elections and higher in low profile elections than it would be if people were more sophisticated and behaved game-theoretically.

REFERENCES

- CAIN, B. E. (1978): "Strategic Voting in Britain," *American Journal of Political Science*, 22(3), 639–655.
- COX, G. W. (1994): "Strategic Voting Equilibria under the Single Nontransferable Vote," *American Political Science Review*, 88(3), 608–621.
- DUVERGER, M. (1954): *Political Parties: Their Organization and Activity in the Modern State*. Wiley, New York.

- FISHER, S. D. (2000): "Intuition versus Formal Theory: Tactical Voting in England 1987–1997," Paper prepared for APSA Annual Meeting, Washington DC.
- KAGEL, J. H., AND A. E. ROTH (1995): *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.
- MORTON, R. H. (2001): *How Elections Rule American Politics*. Norton, forthcoming.
- MYATT, D. P. (2000a): "Strategic Voting Incentives in a Three Party System," Nuffield College Discussion Paper.
- (2000b): "Strategic Voting under the Qualified Majority Rule," Nuffield College Discussion Paper.
- MYERSON, R. B., AND R. J. WEBER (1993): "A Theory of Voting Equilibria," *American Political Science Review*, 87(1), 102–114.
- PALFREY, T. R. (1989): "A Mathematical Proof of Duverger's Law," in *Models of Strategic Choice in Politics*, ed. by P. C. Ordeshook, pp. 69–92. University of Michigan Press, Ann Arbor.
- RIKER, W. H. (1982): *Liberalism Against Populism: A Confrontation Between the Theory of Democracy and the Theory of Social Choice*. Waveland Press, Prospect Heights, IL.