

On the Theory of Strategic Voting¹

DAVID P. MYATT
University of Oxford

First version received March 2003; final version accepted May 2006 (Eds.)

In a plurality-rule election, a group of voters must coordinate behind one of two challengers in order to defeat a disliked status quo. Departing from existing work, the support for each challenger must be inferred from the private observation of informative signals. The unique equilibrium involves limited strategic voting and incomplete coordination. This is driven by negative feedback: an increase in strategic voting by others reduces the incentives for a voter to act strategically. Strategic-voting incentives are lower in relatively marginal elections, after controlling for the distance from contention of a trailing preferred challenger. A calibration applied to the U.K. General Election of 1997 is consistent with the impact of strategic voting and the reported accuracy of voters' understanding of the electoral situation.

1. THE STRATEGIC-VOTING PROBLEM

The winner of a plurality-rule election is the candidate who receives the largest number of votes; there are no prizes for second place. With only two candidates, a voter's decision is simple: she should vote for her favourite. With three (or more) candidates, however, plurality-rule elections are vulnerable to strategic voting: if her preferred candidate is expected to do badly, then a voter might well switch away towards one of the perceived leaders, in the hope that she may exert a greater influence over the outcome of the election.²

The 1970 New York senatorial election provides a classic illustration. In a "three-horse race" two liberal candidates, Richard L. Ottinger and Charles E. Goodell, competed against the conservative James R. Buckley (Table 1). In this scenario the optimal choice for conservative voters was arguably clear: vote for Buckley. Fixing the conservative vote, liberals faced a dilemma. If all had coordinated behind a single challenger, then their combined strength (a total of 3,605,704 votes) would have been sufficient to defeat Buckley. On the other hand, by remaining loyal to their favourite candidates, liberal voters ran the risk of generating a split. Inspecting Table 1, the election's outcome involved only limited coordination of the liberal vote. While there may well have been strategic voting, perhaps from Goodell to Ottinger, it was not enough to prevent Buckley from winning the election.

This depiction of the problem faced by liberal voters generates a stylized *qualified-majority voting* game: a qualified majority (in the New York case a fraction $\frac{38.82}{61.18} \approx 63.5\%$) of a group of voters (the 3,605,704 New York liberals) must coordinate behind one of two candidates (either Ottinger or Goodell) in order to avoid a disliked "status quo" outcome (a win by Buckley). While sharing a dislike for the status quo, the voters differ in their preferences over the challengers. Thus it may be the case that, if the qualified majority is to be attained, some voters

1. Some of the analysis was previously circulated in the papers "A New Theory of Strategic Voting" and "Strategic Voting Under the Qualified Majority Rule".

2. The literature does not distinguish between "strategic" and "tactical" voting; I use "strategic" here. A modern definition (Fisher, 2005) is that "a tactical voter is someone who votes for a party they believe is more likely to win than their preferred party, to best influence who wins in the constituency".

TABLE 1
The 1970 New York Senatorial Election

Candidate	Votes	Share (%)
James R. Buckley	2,288,190	38.82
Charles E. Goodell	1,434,472	24.34
Richard L. Ottinger	2,171,232	36.84
Total	5,893,894	100.00
Conservatives (Buckley)	2,288,190	38.82
Liberals (Goodell + Ottinger)	3,605,704	61.18

Notes: Goodell was a 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, supported instead, the conservative Buckley. The election generated a three-horse race, which has been used as an example of the failed coordination of strategic voting by Riker (1982*a*), Cox (1994), Morton (2006), and others.

will need to vote strategically, by switching towards a second choice. This game captures many elements of more recent plurality elections, as well as the historic New York example. In the 1997 U.K. General Election the incumbent (and unpopular) Conservative party polled between one-third and one-half of the votes for the three major parties in 270 out of 529 English constituencies. In most of England, therefore, anti-Conservative voters needed to successfully coordinate behind either the Labour candidate or the Liberal Democrat candidate in order to ensure a Tory defeat.³

In this paper, I build a tractable model of this qualified-majority voting game. Unlike existing work, there is no common knowledge of the electoral situation. Instead, voters learn about the popularity of the competing candidates, and hence the incentive to vote strategically, via the private observation of informative signals. I identify a unique strategic-voting equilibrium in which voters engage in some, but not complete, strategic voting; hence, voters partially coordinate. Comparative-static exercises permit an assessment of the influence and impact of various factors, including the true underlying popularity of the candidates, the intensity of preferences, and the precision of voters’ information sources.

The prediction that the trailing challenger suffers incomplete strategic desertion is consistent with election outcomes and yet contrasts starkly with existing theories of strategic voting.⁴ Palfrey (1989), Myerson and Weber (1993), and Cox (1994) offered models with “Duvergerian” equilibria in which all votes for the second challenger vanish; for the New York case, this would imply the complete coordination of the liberal vote. The terminology follows from Duverger’s Law (1954), which asserts (Riker, 1982*b*) that “plurality election rules bring about and maintain two-party competition”. There are also “non-Duvergerian” equilibria in which two challengers precisely tie; for instance, an exact 50:50 liberal split between Goodell and Ottinger. The authors cited here assumed that voters have common knowledge of the distribution from which voters’ types (that is, their preferences), and hence their decisions, are independently drawn. In a large electorate, therefore, voters are able to predict almost perfectly the vote shares of the candidates;

3. Throughout the parliamentary constituencies of England, the three major political parties (Conservative, Labour, and Liberal Democrat) are the only effective competitors. In 1997 the Conservative (equivalently, Tory) party was seen by many as right wing, where as the others were seen as left wing. Some might argue that these clear positions of the parties on a Downs–Black–Hotelling spectrum may have changed since Labour’s rise to power. For that reason, the analysis of Section 4.2 restricts attention to the 1997 election.

4. In early decision-theoretic papers (Riker and Ordeshook, 1968; Farquharson, 1969; McKelvey and Ordeshook, 1972; Hoffman, 1982; Cox, 1984; Niemi, 1984) voters do not anticipate strategic switching by others. Later studies (Palfrey, 1989; Myerson and Weber, 1993; Cox, 1994) used game-theoretic models.

this reflects an absence of any real uncertainty in these models. Moreover, the fully coordinated Duvergerian equilibria rely critically upon the assumed common knowledge of the electoral situation.

To see why common knowledge is important, observe that a voter influences the election only when she is *pivotal*. For the New York case, this happens when a liberal candidate exactly ties with the disliked Buckley. In a large electorate, such a tie will almost always involve the leading challenger. If voters commonly understand that Ottinger is the leading challenger, then all will optimally abandon Goodell; the election reduces to a two-horse Duvergerian race. The only way to stop this happening is for the two challengers to run neck and neck. This, in turn, means that a precisely calculated fraction of voters who prefer the more popular challenger must switch *away* from him and towards the less popular challenger.⁵ This somewhat crazy construction supports a non-Duvergerian equilibrium with a critical amount of strategic voting in the *wrong* direction. Putting aside the pathological non-Duvergerian cases, common knowledge of the electoral situation leads to Palfrey's (1989) conclusion that "with instrumentally rational voters and fulfilled expectations, multi-candidate contests under the plurality rules should result in only two candidates getting any votes".

Suppose instead that each voter must use a private signal to infer the identity of the leading challenger. This removal of common knowledge generates a "global game" in the sense of Carlsson and van Damme (1993); it is (Morris and Shin, 2003) a game "of incomplete information whose type space is determined by the players each observing a noisy signal of the underlying state". Here the "underlying state" is the true relative popularity of the challenging candidates, and the "noisy signal" corresponds to voters' information sources, such as the social communication of others' preferences and, indeed, each voter's own preference.

While the removal of common-knowledge assumptions has led to new insights in many settings, little attention has been paid to voting problems.⁶ A notable exception is work by Feddersen and Pesendorfer (1997, 1998), who studied jury models in which each juror's vote is based on a private signal of the defendant's guilt. In their world, there is a single pivotal event, and jurors' preferences (a shared desire to convict the guilty and acquit the innocent) are (at least partially) aligned; a "common value" model. Voters condition on being pivotal in order to update their expected pay-offs from conviction and acquittal. This is related to the winner's curse in common-value auctions, or what Feddersen and Pesendorfer (1996) called a "swing voter's curse". In contrast, this paper offers a "private value" model. A voter knows what she likes, and need not condition on pivotal events in order to work out her preferences. Instead, she uses her private signal to form beliefs about the likely support for the candidates, and thus assesses the relative likelihood of different pivotal events.

Interestingly, the Feddersen–Pesendorfer models give rise to *negative feedback*. Suppose, for instance, that a conviction requires unanimity. If her colleagues are biased towards a guilty verdict, then a juror will realize that her own signal is the only source of real information, and will truthfully base her vote on it. Equivalently, the bias of others feeds back to a reduction in an individual's bias. Conversely, when others are expected to "vote their signals" then a juror will realize that being pivotal reveals strong evidence of the defendant's guilt; she will then ignore her own signal and always vote for a conviction. This story ensures that an equilibrium must involve some, but not complete, bias in voting decisions.

Negative feedback is also present in this paper. When others vote strategically by responding strongly to their private signals (for instance, a Goodell supporter will switch if she perceives

5. Here the "leading challenger" is the candidate who, given voting behaviour, is in a better position to win, whereas the "most popular challenger" is the candidate who most anti-Buckley voters would prefer to win.

6. The techniques of global games have been applied to problems including currency crises (Morris and Shin, 1998), bank runs (Goldstein and Pauzner, 2005), and debt pricing (Morris and Shin, 2004).

only a tiny bias in popularity towards Ottinger) then an individual will tend to stick with her first preference. To understand why, notice that, when others follow their signals, only a small lead in underlying popularity is needed for a challenger to hit the magical qualified majority. Thus, the pivotal events (successful challenges by Goodell and Ottinger, respectively) arise from underlying electoral situations (that is, relative popularities) that are very close, and hence have similar probabilities. Since the likelihood ratio of these pivotal events will be close to 1, the incentive to vote strategically will tend to be small.

This negative feedback contrasts with the positive-feedback story of the “bandwagon” effect (Simon, 1954). The bandwagon hypothesis is that the loss of support from a less popular challenger enhances the incentive to vote strategically, hence further eroding the support of the trailing challenger. The bandwagon story cannot apply here: when voters do not have a common understanding of the electoral situation, it will not be clear who the leading challenger is. Informally, a voter becomes fearful that the bandwagon is rolling in a direction opposite to that indicated by her own private signal, and is more cautious about switching her vote. Anticipation of switching by others (a speedier bandwagon) amplifies this caution.

Various implications flow from the analysis. First, negative feedback ensures that strategic voting is lower in equilibrium than it would be if each strategic voter naively expected others to vote sincerely. Second, since private signals sometimes point the wrong way, strategic voting will be bidirectional. The net effect, however, will be to bolster the support of the leading challenger. Third, a unique equilibrium with only partial coordination permits interesting comparative-static exercises: from these emerge the implications of the electoral situation; of the intensity of preferences; and of the quality of information sources. Finally, a calibration of the model offers insights into the wider consequences of strategic voting.

In Section 2 a qualified-majority voting game captures the “beat the disliked conservative” scenario and yields a unique strategic-voting equilibrium. In Section 3 I interpret the results in the context of plurality-rule elections. In Section 4 calibrations of the model reveal the potential impact of strategic voting. Formal proofs are relegated to Appendix A.

2. STRATEGIC-VOTING EQUILIBRIA

I have suggested that New York liberals were playing a game of “beat the disliked conservative”. Here the stylized features are captured by a model of qualified-majority voting.

2.1. *A model of qualified-majority voting*

In this section I describe the players, their preferences, the voting rules, and the information upon which decisions are based.

2.1.1. Voting rules. The $n + 1$ members of an electorate participate in a simultaneous-move binary-action game. Each player $i \in \{0, 1, \dots, n\}$ must vote for one of two candidates $j \in \{1, 2\}$, yielding vote totals of x_1 and x_2 , which satisfy $x_1 + x_2 = n + 1$. Candidate j wins the election if and only if he captures (strictly) more than \bar{x} votes, where $\frac{n+1}{2} \leq \bar{x} \leq n$. If neither candidate wins, then the outcome is the status quo. Notice that if $\bar{x} = \frac{n+1}{2}$ then a candidate needs only a simple majority in order to win.⁷ In contrast, when \bar{x} is larger, a “qualified majority” of players need to coordinate their votes on a single candidate if the status quo is to be overturned. The fraction $\frac{\bar{x}}{n+1}$ is the degree of coordination that is required.

7. When $\bar{x} = \frac{n+1}{2}$ and $(n + 1)$ is even, then a 50:50 tie $x_1 = x_2 = \bar{x}$ is possible. According to the rules given here, the status quo would prevail. This specification may be changed without loss of generality: nothing of interest would be affected if, in this case, a winner $j \in \{1, 2\}$ were chosen via the toss of a fair coin.

2.1.2. Pay-offs. Voter i receives a zero pay-off from the status quo, and a positive pay-off $u_{ij} > 0$ for a win by candidate j . Subsequent analysis confirms that the ratio of u_{i1} and u_{i2} determines optimal voting behaviour. Writing $\tilde{u}_i \equiv \log \left[\frac{u_{i1}}{u_{i2}} \right]$, the sign of \tilde{u}_i reveals the identity of voter i 's preferred candidate, and $|\tilde{u}_i|$ is the intensity of her preference.⁸ I assume that this log relative preference \tilde{u}_i satisfies $\tilde{u}_i = \eta + \varepsilon_i$. Here, η is a *common component* to the preferences of all voters; it is an underlying state variable for the whole game. Conditional on η , the *idiosyncratic component* is an independent draw for each voter, and satisfies $\varepsilon_i \mid \eta \sim N(0, \xi^2)$.

Under this specification, η captures the preferences of a median voter. Furthermore, the probability $\pi \equiv \Pr[\tilde{u}_i > 0 \mid \eta]$ that a voter's favourite challenger is the first candidate satisfies $\pi = \Phi(\eta/\xi)$, where $\Phi(\cdot)$ is the distribution function of the standard normal. Thus η indexes the relative support of the candidates, and the first candidate is more popular if and only if $\eta > 0$. (It is equivalent to specify π as the state variable and then generate $\eta = \xi \Phi^{-1}(\pi)$.) The variance $\xi^2 = \text{var}[\tilde{u}_i \mid \eta]$ measures the idiosyncrasy of voters' preference intensities.

2.1.3. Information. Voter i knows her own preferences, but not those of others. Furthermore, she does not know the decomposition of \tilde{u}_i into its common and idiosyncratic components. The ignorance of η is reflected in an improper prior over η .⁹ Following the observation of \tilde{u}_i , player i updates to obtain the posterior belief $\eta \mid \tilde{u}_i \sim N(\hat{\eta}_i, \kappa^2)$: absent any other information, a voter assesses the true underlying support for the candidates based on her own preferences.

Fortunately, voter i has access to a second information source: a signal s_i where conditional on η , $s_i \sim N(\eta, \sigma^2)$ independently across the electorate. To allow for correlation between this signal and a voter's preferences, I assume that s_i and \tilde{u}_i are joint normal (conditional on η) with correlation coefficient ρ . A voter's type, therefore, is the preference-signal pair (\tilde{u}_i, s_i) . Based on this type, a voter forms updated beliefs about η .

Lemma 1. *A voter's updated beliefs satisfy $\eta \mid (\tilde{u}_i, s_i) \sim N(\hat{\eta}_i, \kappa^2)$, where*

$$\hat{\eta}_i = ws_i + (1 - w)\tilde{u}_i, \quad w = \frac{\xi^2 - \rho\xi\sigma}{\xi^2 + \sigma^2 - 2\rho\xi\sigma}, \quad \text{and } \kappa^2 = w^2\sigma^2 + (1 - w)^2\xi^2 + 2w(1 - w)\rho\xi\sigma.$$

Conditional on the state, $\hat{\eta}_i \mid \eta \sim N(\eta, \kappa^2)$, $\text{cov}[\tilde{u}_i, \hat{\eta}_i \mid \eta] = \kappa^2$, and κ^2 satisfies $\kappa^2 \leq \xi^2$.

The posterior expectation $\hat{\eta}_i$ is a sufficient statistic for inference about η ; it fully captures a voter's beliefs about the underlying support for the candidates. Leaning upon this, it is convenient to redefine a player's type as the pair $(\tilde{u}_i, \hat{\eta}_i)$, which represents preferences and beliefs about the state η . Henceforth, all reference to the signal s_i will be dropped, and I work directly with $(\tilde{u}_i, \hat{\eta}_i)$. I assume (with little loss of generality) that $\kappa^2 < \xi^2$.

2.1.4. Commentary. The shared distaste for the status quo ensures that a qualified-majority election resembles an $(n + 1)$ -player battle-of-the-sexes: when $u_{i1} > u_{i2}$ voter i would prefer to see a win by the first candidate, but would be willing to switch if she believed that her vote would nudge the second candidate's vote count over the qualified-majority threshold. The game captures, therefore, the stylized features of strategic voting under the plurality rule.

8. To obtain the actual pay-offs u_{i1} and u_{i2} from \tilde{u}_i , simply set $u_{i1} = e^{\tilde{u}_i}/[1 + e^{\tilde{u}_i}]$ and $u_{i2} = 1 - u_{i1}$.

9. Beginning with a common normal prior over η , the strategic-voting equilibrium obtained converges to the one obtained with an improper prior as the variance of the prior diverges to ∞ (see Section 2.3).

2.2. Optimal behaviour and voting equilibria

The parameters ξ^2 (variance of preference intensity) and κ^2 (variance of voters' expectations) are common knowledge, but the state variable η is not. Here I investigate how a voter's beliefs about the electoral situation (captured by $\hat{\eta}_i$) interact with her preferences as part of a strategic-voting equilibrium in a large electorate. As there will be a focus on expanding electorate sizes ($n \rightarrow \infty$), the qualified-majority hurdle depends on n and is assumed to satisfy $\frac{\bar{x}}{n} \rightarrow \gamma \in (\frac{1}{2}, 1)$ as $n \rightarrow \infty$. To ease notation, and without loss of generality, \bar{x} will take only integer values.

2.2.1. Voting strategies. Voter i 's action, characterized by the probability $v_i \in [0, 1]$ that she votes for the first candidate, will depend upon her preference-belief type $(\tilde{u}_i, \hat{\eta}_i)$. A *voting strategy* for i is then a mapping from a type to an action $v_i(\tilde{u}_i, \hat{\eta}_i) : \mathbb{R}^2 \mapsto [0, 1]$. Given her pay-offs, she votes *sincerely* if she votes for her favourite candidate. Hence, if $u_{i1} > u_{i2}$, or equivalently $\tilde{u}_i > 0$, a sincere vote would be for the first candidate; on the other hand, if she finds it optimal to vote for the second candidate, then she votes *strategically*.

Absent further restrictions, there are many Nash equilibria. For instance, if voters ignore their types and all vote for the same candidate, then they will never be pivotal and hence will be indifferent between their actions. Similar arguments reveal that any type-independent pure-strategy profile yielding vote totals (x_1, x_2) is an equilibrium, unless $x_j + 1 > \bar{x} \geq x_j$ for some $j \in \{1, 2\}$. (For the exceptions, everyone would switch to candidate j .)

Type-independent equilibria have undesirable features: voters ignore what they know about themselves and the electoral situation, they need to follow individual-specific actions, and such equilibria are not robust to non-instrumental pay-off considerations. Exploring this last critique, note when a voter is indifferent between her actions, she might vote sincerely to express support for her favourite candidate. In fact, if a lexicographic preference structure is imposed, so that a voter acts sincerely when indifferent, type-independent strategy profiles are no longer equilibria.¹⁰ In response to these observations, attention turns to voting strategies that respond positively to private information but not to the labelling of the players.

Definition 1. A strategy profile is *symmetric* if the same voting strategy $v(\tilde{u}_i, \hat{\eta}_i)$ is used by each player. It exhibits *multi-candidate support* if there are types that yield votes of $v(\tilde{u}_i, \hat{\eta}_i) = 0$ and $v(\tilde{u}_i, \hat{\eta}_i) = 1$. It is *monotonic* if $v(\cdot)$ is (weakly) increasing in its arguments.

Attention is restricted to strategies satisfying these criteria. The first two properties respond to the critiques presented above. They ensure that all outcomes occur with positive probability, and hence voters can envisage being pivotal. This means that, even when expressive concerns are present, instrumental motivations will drive behaviour. The third property requires a positive response to the perceived popularity of a candidate.

2.2.2. Voting in large electorates. Given that all voters $i \in \{1, \dots, n\}$ adopt the same voting strategy $v(\tilde{u}_i, \hat{\eta}_i)$, consider the pay-offs of voter 0. Dropping her identity subscript for

10. At first blush, this critique might seem to apply to type-dependent equilibria: in a large electorate the absolute probability of a pivotal outcome is vanishingly small, and so a voter will be almost indifferent between the candidates. The refinement of "vote sincerely when indifferent" would then apply. This position assumes, however, that a voter's pay-off u_{ij} for a win by candidate j is constant with respect to the electorate size. An alternative specification would be one in which the pay-off grows in proportion to the electorate, so that a voter cares more about election results that affect more people. As n grows the absolute probability of a pivotal outcome falls, but the degree to which a voter cares about the outcome rises in tandem. Hence, even for large n , a voter's instrumental concerns remain central to her voting decision.

simplicity, her expected pay-offs from actions $j \in \{1, 2\}$, as a function of her type, are $U_1(\tilde{u}, \hat{\eta} | v)$ and $U_2(\tilde{u}, \hat{\eta} | v)$. It is straightforward to confirm that there is zero probability that the voter will be indifferent between her actions, and so it is without loss of generality to suppose that she votes for the first candidate when indifferent. Writing $\mathbb{I}[\cdot]$ for the indicator function, her best response is simply $\text{BR}(\tilde{u}, \hat{\eta} | v) \equiv \mathbb{I}[U_1(\tilde{u}, \hat{\eta} | v) \geq U_2(\tilde{u}, \hat{\eta} | v)]$.

A voting strategy is a Nash equilibrium if it satisfies $v(\tilde{u}, \hat{\eta}) = \text{BR}(\tilde{u}, \hat{\eta} | v)$. A Nash solution will depend upon the electorate size $n + 1$ and the qualified majority \bar{x} . My aim, however, is to characterize equilibria in large electorates: a voting strategy that works for large n .

One possibility would be to find Nash voting strategies (if they exist) for each n , and examine their properties (assuming their sequence converges) as $n \rightarrow \infty$. I take a slightly different approach, by defining a solution concept directly over a sequence of voting games.

Definition 2. A *strategic-voting equilibrium* is a symmetric and monotonic strategy $v^*(\tilde{u}, \hat{\eta})$ exhibiting multi-candidate support, such that $\Pr[v^*(\tilde{u}, \hat{\eta}) = \lim_{n \rightarrow \infty} \text{BR}(\tilde{u}, \hat{\eta} | v^*)] = 1$.

Heuristically, this defines an “ ε -equilibrium” for large electorates and arbitrarily small ε . Hence a strategic-voting equilibrium is a voting strategy such that, when it is adopted, leads to the play of a best response by (almost) all types in electorates that are sufficiently large.¹¹

In one sense, a strategic-voting equilibrium is less stringent than the usual Nash solution concept. Given that $v^*(\tilde{u}, \hat{\eta})$ is played, there may be types who do not play a best response for a particular finite n . In a second sense, a strategic-voting equilibrium is more stringent than Nash. Whereas a Nash equilibrium would involve the play of a best response by almost all types for a particular finite n , it would not necessarily be robust to increases in n .¹²

2.2.3. Best responses and linear voting strategies. To characterize strategic-voting equilibria, I begin by calculating a voter’s best response $\text{BR}(\tilde{u}, \hat{\eta} | v)$. Focusing on voter 0, x is the number of votes cast for the first candidate by the remaining n -strong electorate. Voter 0 recognizes that her vote counts only when she is pivotal: a vote for the first candidate changes her pay-off from the zero-pay-off status quo to u_1 if and only if $x = \bar{x}$, and, similarly, a vote for the second candidate yields a gain of u_2 if and only if $n - x = \bar{x}$. Hence voter 0 should vote for the first candidate if and only if $u_1 \Pr[x = \bar{x} | \hat{\eta}] \geq u_2 \Pr[x = n - \bar{x} | \hat{\eta}]$. On re-arranging,

$$\text{BR}(\tilde{u}, \hat{\eta} | v) = \mathbb{I} \left[\tilde{u} + \log \left[\frac{\Pr[x = \bar{x} | \hat{\eta}]}{\Pr[x = n - \bar{x} | \hat{\eta}]} \right] \geq 0 \right]. \tag{1}$$

In acting optimally, a voter balances her relative preference for two competing candidates, as captured by \tilde{u} , against the relative likelihood of pivotal events. This log-likelihood ratio, which depends upon the voter’s beliefs $\hat{\eta}$, represents her incentive to vote strategically.

To evaluate (1), a voter must form beliefs over x . Conditional on the state η , the actions of others are independent: $v_i = 1$ with probability $p = E[v(\tilde{u}_i, \hat{\eta}_i) | \eta]$, and hence x is drawn from a binomial distribution with parameters p and n . As voter 0 is ignorant of η and hence p , she must take expectations over p , conditional on the sufficient statistic $\hat{\eta}$.

11. To expand on this description, consider $(\tilde{u}, \hat{\eta})$ where $\text{BR}(\tilde{u}, \hat{\eta} | v) \rightarrow v(\tilde{u}, \hat{\eta}_v)$. Since $\text{BR}(\tilde{u}, \hat{\eta} | v) \in \{0, 1\}$ it must be that $\text{BR}(\tilde{u}, \hat{\eta} | v) = v(\tilde{u}, \hat{\eta})$ for all n large enough. When evaluating $\Pr[v^*(\tilde{u}, \hat{\eta}) = \lim_{n \rightarrow \infty} \text{BR}(\tilde{u}, \hat{\eta} | v^*)]$, any continuous distribution over $(\tilde{u}, \hat{\eta})$ can be used; for instance, that arising from any particular state η .

12. The existence of Nash equilibria with multi-candidate support for a particular electorate size n is established and discussed in Appendix B.2.

Lemma 2. Fix a symmetric and monotonic voting strategy $v(\tilde{u}_i, \hat{\eta}_i)$ exhibiting multi-candidate support. Then the distribution $F(p \mid \hat{\eta}, v) \equiv \Pr[E[v(\tilde{u}_i, \hat{\eta}_i) \mid \eta] \leq p \mid \hat{\eta}]$ admits a continuous density function $f(p \mid \hat{\eta}, v)$, which is strictly positive for $p \in (0, 1)$.

Given this, voter 0 can calculate the probabilities of the two pivotal events. For instance,

$$\Pr[x = \bar{x} \mid \hat{\eta}] = \int_0^1 \binom{n}{\bar{x}} p^{\bar{x}} (1-p)^{n-\bar{x}} f(p \mid \hat{\eta}, v) dp > 0, \tag{2}$$

with a similar expression for $\Pr[x = n - \bar{x} \mid \hat{\eta}] > 0$. These probabilities vanish to 0 as n grows large, but their ratio is well behaved. As a consequence of the Law of Large Numbers, $n \times \Pr[x = \bar{x} \mid \hat{\eta}] \rightarrow f(\gamma \mid \hat{\eta}, v)$ as $n \rightarrow \infty$ (Chamberlain and Rothschild, 1981; Proposition 1). Heuristically, vote shares will converge to p and $1 - p$ as $n \rightarrow \infty$, so that a voter is pivotal only when p is close to either γ or $1 - \gamma$. In fact,

$$\lim_{n \rightarrow \infty} \left[\frac{\Pr[x = \bar{x} \mid \hat{\eta}]}{\Pr[x = n - \bar{x} \mid \hat{\eta}]} \right] = \frac{f(\gamma \mid \hat{\eta}, v)}{f(1 - \gamma \mid \hat{\eta}, v)}. \tag{3}$$

Thus a voter computes the likelihood ratio that the first candidate’s support p just hits two critical values. The normality of voters’ beliefs over the underlying state variable η ensures that the log-likelihood ratio of pivotal events takes a linear form as $n \rightarrow \infty$.

Lemma 3. Fix a symmetric and monotonic voting strategy $v(\tilde{u}_i, \hat{\eta}_i)$ exhibiting multi-candidate support. A voter’s best response satisfies $\text{BR}(\tilde{u}, \hat{\eta} \mid v) \rightarrow \mathbb{I}[\tilde{u} + \lambda(\hat{\eta}) \geq 0]$ as $n \rightarrow \infty$, where $\mathbb{I}[\cdot]$ is the indicator function, and where

$$\lambda(\hat{\eta}) \equiv \log \left[\frac{f(\gamma \mid \hat{\eta}, v)}{f(1 - \gamma \mid \hat{\eta}, v)} \right] = a + b\hat{\eta} \text{ where } a \in \mathbb{R} \text{ and } b \in \mathbb{R}_+. \tag{4}$$


Fixing a linear voting strategy $v(\tilde{u}_i, \hat{\eta}_i) = \mathbb{I}[\tilde{u}_i + a + b\hat{\eta}_i \geq 0]$ for some $a \in \mathbb{R}$ and $b \in \mathbb{R}_+$, a voter’s best response satisfies $\text{BR}(\tilde{u}, \hat{\eta} \mid v) \rightarrow \mathbb{I}[\tilde{u} + \hat{a} + \hat{b}\hat{\eta} \geq 0]$ as $n \rightarrow \infty$, where \hat{a} and \hat{b} satisfy

$$\hat{a} = \hat{a}(a, b) \equiv \frac{a\hat{b}(b)}{(1+b)}, \quad \text{and} \quad \hat{b} = \hat{b}(b) \equiv \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + (b^2 + 2b)\kappa^2}}{\kappa^2(1+b)}. \tag{5}$$

When voters use the same monotonic strategy exhibiting multi-candidate support, the probability $p = E[v(\tilde{u}_i, \hat{\eta}_i) \mid \eta]$ that i votes for the first candidate is a strictly increasing and continuous function of the underlying state η . Hence, there will be a critical state η_1 at which the support of the first candidate just equals the qualified majority γ and a second critical state η_2 where the second candidate does the same. $\lambda(\hat{\eta})$ from Lemma 3 will be proportional to the log-likelihood ratio of η_1 and η_2 , conditional on a voter’s beliefs $\hat{\eta}$. Given the normality assumption, $\lambda(\hat{\eta})$ will be linear in the posterior mean $\hat{\eta}$, so that $\lambda(\hat{\eta}) = a + b\hat{\eta}$.

Linear voting strategies are easy to interpret. The intercept a is the strategic incentive faced by a voter given that she has neutral beliefs $\hat{\eta} = 0$. It is a single-independent bias towards candidate 1 (when $a > 0$) or candidate 2 (when $a < 0$). Since $\hat{a} = \hat{b}a/(1+b)$ and $b > 0$, it is clear that a voter will tend to exhibit a bias (i.e. $\hat{a} > 0$) if and only if she expects others to exhibit such a bias (when $a > 0$). In contrast, b represents the response of the strategic incentive to a voter’s private beliefs about the likely electoral situation.

For $v(\tilde{u}, \hat{\eta})$ to yield a strategic-voting equilibrium, $v(\tilde{u}, \hat{\eta}) = \lim_{n \rightarrow \infty} \text{BR}(\tilde{u}, \hat{\eta} \mid v)$. Lemma 3 demonstrates that the only strategies satisfying this criterion are those that are linear.



Notes: To illustrate negative feedback, consider first sincere voting. The first candidate wins when $\eta > \eta^\dagger > 0$ where $\gamma = \Pr[\tilde{u}_i > 0 \mid \eta^\dagger]$; similarly the second candidate wins when $\eta < -\eta^\dagger$. A voter evaluates the relative likelihood of η^\dagger vs. $-\eta^\dagger$. For the beliefs illustrated, η^\dagger is more likely so there is a strategic incentive to vote for the first candidate. Consider next a situation in which voters respond to their signals, so that $b > 0$. A voter will now compare the relative likelihood of η^\ddagger and $-\eta^\ddagger$ where $\gamma = \Pr[\tilde{u}_i + b\hat{\eta}_i \geq 0 \mid \eta^\ddagger]$. Now, when $b > 0$, $0 < \eta^\ddagger < \eta^\dagger$; it takes a smaller lead (in terms of η) for the first candidate to be in a knife-edge position; similarly for the second candidate. By inspection, the odds of η^\ddagger vs. $-\eta^\ddagger$ are smaller than the odds of η^\dagger vs. $-\eta^\dagger$; the strategic-voting incentive falls.

FIGURE 1

The negative-feedback effect

2.2.4. Negative feedback and voting equilibria. Strategic voting may involve negative feedback. To verify this, consider the slope b of a linear voting strategy. When b is high, voters respond strongly to their beliefs. This increases strategic voting in *both* directions. When $\hat{\eta} > 0$, a voter worries that others may hold beliefs $\hat{\eta}_i < 0$, which might lead to a pivotal event involving candidate 2 rather than candidate 1. For high $\hat{\eta}$, this seems unlikely; surely the first candidate will almost certainly win? But if candidate 1 will almost certainly win, then a single vote has no effect. In fact, if a tie occurs then it suggests that the signal $\hat{\eta}$ was wildly optimistic. A voter must envisage a true value of η much lower than that indicated by her signal. Of course, this leads her to contemplate state variables satisfying $\eta < 0$. Figure 1 provides a graphical illustration of the negative-feedback effect; an increase in b forces a voter to compute the relative likelihood of underlying states (that is, η) that are closer together.

Formally, the negative-feedback effect is reflected in the fact that the mapping $\hat{b}(b)$ from (5) in Lemma 3 is decreasing in b . This ensures the existence of a unique fixed point $b^* > 0$. Fixing b^* , the unique fixed point of the mapping $\hat{a}(a, b^*)$ is $a^* = 0$. Summarizing

Proposition 1. *There is a unique strategic-voting equilibrium $v^*(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + b^*\hat{\eta} \geq 0]$ for some $b^* > 0$, where b^* is increasing in γ , but decreasing in κ^2 and ζ^2 .*

That $a^* = 0$ in equilibrium means that there is no systematic bias towards one candidate. Instead, strategic voting is driven by the voters’ responses to signals of the electoral situation. Of course, signals can be wrong, and so some voters will hold beliefs $\hat{\eta}_i < 0$ even though $\eta > 0$; this yields a strategic incentive to switch away from the more popular candidate. Thus strategic voting is bidirectional, with some voters switching in the wrong direction.

Proposition 1 offers comparative-static predictions: a voter responds more strongly to her beliefs about the underlying state when more coordination is required (γ is higher) and when her beliefs are more precise (κ^2 is lower). (The effect of ζ^2 may be misleading, however, since (fixing κ^2) a change in ζ^2 alters the probability that a voter is able to identify correctly the relative positions of the candidates, and hence the accuracy of her beliefs.)

I postpone further comparative-static exercises until Section 3. Here, attention is turned to a benchmarking exercise. The equilibrium voting strategy takes into account strategic switching by others. It may be compared to a “decision theoretic” voting strategy; a strategy employed by a voter who expects others to vote sincerely.

Proposition 2. *If the electorate adopt a sincere voting strategy $v(\tilde{u}_i, \hat{\eta}_i) = \mathbb{I}[\tilde{u}_i \geq 0]$ then a voter’s best response satisfies $\lim_{n \rightarrow \infty} \text{BR}(\tilde{u}, \hat{\eta} \mid v) = v^\dagger(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + b^\dagger \hat{\eta} \geq 0]$ for some b^\dagger , where $b^\dagger > b^* > 0$ and b^\dagger is increasing in γ but decreasing in κ^2 and ζ^2 .*

The inequality $b^* < b^\dagger$ is a consequence of the negative-feedback effect: it says that game-theoretic considerations (anticipating strategic voting by others) result in a reduced, rather than enhanced strategic response to a voter’s perception of the electoral situation.

2.3. Discussion

Before Section 3’s study of the practical implications and interpretations of the qualified-majority voting model, a selection of the theoretical issues are addressed.

2.3.1. Strategic uncertainty. To calculate an optimal action, a voter envisages different pivotal events, and hence she must be uncertain of the election outcome. Equation (2) reveals two sources of uncertainty. First, each vote is drawn from a Bernoulli distribution with parameter p . This is *idiosyncratic* uncertainty, due to noisy type realizations. Second, voters do not know η , and hence do not know $p = E[v(\tilde{u}_i, \hat{\eta}_i) \mid \eta]$. This is *strategic* uncertainty and is captured by the density $f(p \mid \hat{\eta}, v)$. Equation (3) reveals that, in a large electorate, incentives are driven entirely by strategic uncertainty. When strategic uncertainty is omitted (for instance, when the electoral situation is common knowledge) results are driven by idiosyncratic uncertainty. However, in a large electorate idiosyncrasies are averaged out. This suggests strategic uncertainty is crucial to a useful model of strategic voting.

2.3.2. Public signals. Whereas omitting strategic uncertainty (via common knowledge of the voters’ type distribution) is a strong assumption, the specification here pushes towards an opposite extreme: there is no public signal of the underlying state. To remedy this, suppose that voters commonly observe a public signal $\mu \mid \eta \sim N(\eta, \chi^2)$. (This remains a step away from complete knowledge of the electoral situation, since μ is only a noisy signal of η .) Incorporating the public signal, voter i updates to form beliefs

$$\eta \mid \{\hat{\eta}_i, \mu\} \sim N\left(\frac{\chi^2 \hat{\eta}_i + \kappa^2 \mu}{\chi^2 + \kappa^2}, \frac{\chi^2 \kappa^2}{\chi^2 + \kappa^2}\right). \tag{6}$$

The public signal μ can spark a positive-feedback bandwagon. To see this note firstly that Lemma 2 holds in the presence of the public signal. Secondly Lemma 3 also holds so long as the mapping $\hat{a}(a, b)$ is modified to

$$\hat{a}(a, b) = \frac{\hat{b}(b) \times \mu}{\zeta} + \left[\frac{1 + \zeta}{\zeta} \times \frac{\hat{b}(b)}{1 + b} \times a \right] \quad \text{where} \quad \zeta \equiv \frac{\chi^2}{\kappa^2}. \tag{7}$$

Here ζ indexes the relative importance of public and private information. If others vote sincerely ($a = b = 0$) a voter will bias towards the publicly popular candidate ($\hat{a}(0, 0) = \frac{\mu}{1 + \zeta} > 0$ when $\mu > 0$). Furthermore, a bias on the part of others feeds back into a voter’s best response. In fact, this positive feedback is explosive if $\zeta < b^*$.

Proposition 3. *Suppose that all voters commonly observe a public signal $\mu \mid \eta \sim N(\eta, \chi^2)$. If $\zeta > b^*$, there is a unique strategic-voting equilibrium $v^*(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + a^* + b^*\hat{\eta} \geq 0]$, where*

$$a^* = \frac{b^*(1 + b^*)}{\zeta - b^*} \times \mu. \tag{8}$$

The bias a^ is decreasing in χ^2 . Furthermore, $\lim_{\chi^2 \uparrow \infty} a^* = 0$ and $\lim_{\chi^2 \downarrow b^*\kappa^2} a^* = \pm\infty$.*

Hence the public signal has no effect on a voter’s reaction b^* to her private signal. Instead, it generates a signal-independent bias a^* towards the candidate who enjoys more public popularity. The effect can be dramatic. When the public signal favours the first candidate, so that $\mu > 0$, the bias explodes ($a^* \uparrow \infty$) as $\zeta \downarrow b^*$; the electorate coordinates behind the first candidate, even though the noise in the public signal remains bounded away from 0.

Proposition 3 imposes the restriction $\zeta > b^*$, or equivalently $\frac{\chi^2}{\kappa^2} > b^*$: the variance of the public signal needs to be large relative to the variance of private signals. Note, however, that b^* also depends on κ^2 . In fact, straightforward algebra reveals $\zeta > b^*$ necessitates $\chi^2 > 2\Phi^{-1}(\gamma)\kappa$. This comparison involves the variance of the public signal, but the standard deviation of private signals.¹³ Heuristically, strategic-voting equilibria require an environment in which private signals are much more important than public signals.

So what happens when $\zeta < b^*$? The equilibrium described in Proposition 3 leads to $a^* < 0$ for $\mu > 0$; voters actively move against the publicly popular candidate. This pathological feature is shared by the non-Duvergerian equilibria discussed in Section 1. Such equilibria do not provide robust and convincing predictions of play. In fact, building an appropriate strategy-revision process (such as iteratively updated best response) play converges towards a strategic-voting equilibrium only if $\zeta > b^*$.¹⁴ For $\zeta < b^*$, such a process tends to diverge away towards a fully coordinated strategy profile.

This discussion leads to the following conclusion: when voters observe a sufficiently precise public signal of the electoral situation, then the qualified-majority voting model mimics the plurality-voting models of Palfrey (1989) and Myerson and Weber (1993). It is the common knowledge of the public signal that pushes voters towards complete coordination.

2.3.3. Monotonicity. Returning to purely private information sources, it is easy to confirm that any best response must be increasing in \tilde{u}_i . However, it is possible to construct strategies where voters respond negatively to their beliefs about the electorate situation. Inspecting equation (5), the mapping $\hat{b}(b)$ has a second fixed point satisfying $b^* < -1$. As in the pathological equilibrium of the public-signal model discussed above, voters move against more popular candidates. In fact, candidate 2 almost always wins when $\eta \rightarrow \infty$.

Such equilibria are not robust. To see why, consider the following change: suppose that when $|\tilde{u}_i|$ is sufficiently large, the pay-off from voter i ’s second choice drops to 0, so that she has a weakly dominant strategy to vote for her first choice. This means that, for $\eta \rightarrow \infty$, almost everyone would vote for candidate 1. More generally, this modification would also knock out fully coordinated type-independent equilibria (see Appendix B.1).

13. A similar inequality was derived by Morris and Shin (2004). They studied a coordination game played by creditors who are able to either foreclose on a loan or roll it over. Their “global game” has a unique equilibrium when $\frac{\alpha}{\sqrt{\beta}} < \frac{\sqrt{2\pi}}{z}$, where α is the precision of a public signal, β is the precision of private signals, π is the well-known constant, and z is a parameter indexing the penalties of mis-coordination in their game.

14. Index a sequence of symmetric and monotonic strategies by t , where $v_{t+1}(\tilde{u}, \hat{\eta}) \equiv \lim_{n \rightarrow \infty} \text{BR}(\tilde{u}, \hat{\eta} \mid v_t)$. Suitably modified to incorporate a public signal, Lemma 3 ensures that $v_t(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + a_t + b_t\hat{\eta} \geq 0]$ for $t \geq 1$, where $b_{t+1} = \hat{b}(b_t)$ and $a_{t+1} = \hat{a}(a_t, b_t)$. The properties of $\hat{b}(b)$ ensure that $\lim_{t \rightarrow \infty} b_t = b^*$. However, $\lim_{t \rightarrow \infty} a_t = a^*$ only if $\zeta > b^*$. If $\zeta < b^*$ then $a_t \rightarrow \pm\infty$ as $t \rightarrow \infty$. (See Appendix B.3 for more details.)

3. INTERPRETATION

Here I interpret and explore the qualified-majority voting model described in Section 2. Firstly, I consider possible sources of voting information. Secondly, I compute the determinants of strategic-voting incentives. Finally, I assess the likely impact of strategic voting.

3.1. Information

I have already described one information source: a voter's observation of \tilde{u}_i . Before discussing other sources, I consider the accuracy of voters' beliefs.

3.1.1. Accuracy. So far, the noise in voters' beliefs has been indexed by κ^2 . Here I write α for probability that a voter i correctly identifies the leading candidate; α is the *accuracy* of her beliefs. To illustrate, suppose that the first candidate is more popular, so that $\eta > 0$. Voter i correctly identifies the leading candidate whenever $\hat{\eta}_i > 0$. Recalling that $\hat{\eta}_i | \eta \sim N(\eta, \kappa^2)$, the accuracy of her beliefs is simply $\alpha \equiv \Pr[\hat{\eta}_i > 0 | \eta] = \Phi(\eta/\kappa)$. Now, recall from Section 2.1 that the probability that voter prefers candidate 1 (representing candidate 1's true underlying support) is $\pi = \Phi(\eta/\xi)$, and so $\eta = \xi \Phi^{-1}(\pi)$. Substituting back in to the accuracy measure,

$$\alpha = \Phi\left(\frac{\xi}{\kappa} \times \Phi^{-1}(\pi)\right), \quad \text{or equivalently} \quad \kappa = \xi \times \frac{\Phi^{-1}(\pi)}{\Phi^{-1}(\alpha)}. \quad (9)$$

By inspection, α is jointly determined by κ^2 and the idiosyncrasy parameter ξ^2 .

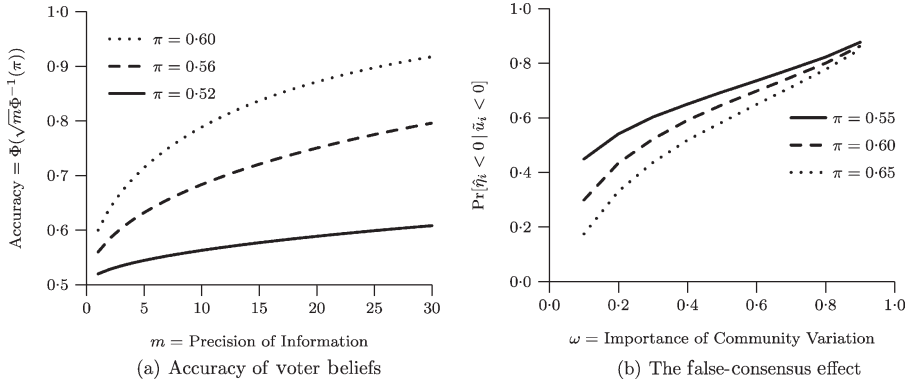
3.1.2. Social communication. Moving beyond introspection, voters learn via communication with others.¹⁵ To formalize this, suppose that voter i observes m independent draws from $N(\eta, \xi^2)$, including the realization of \tilde{u}_i . The sample mean is a sufficient statistic for inference about η , and forms the mean of her posterior beliefs $\eta | \hat{\eta}_i \sim N(\hat{\eta}_i, \kappa^2)$ where $\kappa^2 = \frac{\xi^2}{m}$. Thus, beliefs with precision m/ξ^2 (the reciprocal of the variance) arise from a "private opinion poll" of m voters.¹⁶ For $\pi > \frac{1}{2}$, the accuracy of beliefs is $\alpha = \Phi(\sqrt{m} \times \Phi^{-1}(\pi))$. Hence the understanding of the electoral situation will improve with the size of social networks (Figure 2(a)).

A voter is not always limited to social communication. A further information source might be opinion polls or other media input. Opinion polls are rare at the level of a constituency.¹⁷ Nevertheless, a poll may be viewed within the context of the private-signal specification (cf. Section 2.3). If a poll perfectly identifies η , then κ^2 would correspond to any (possibly small) noise in a voter's observation of it. Based on this, α can be interpreted as the accuracy of a voter's observation of the media. As an illustration, setting $\alpha = 0.95$ coupled with $\pi = 0.6$ would be

15. Pattie and Johnston (1999) demonstrated that the contextual effects of conversations with family, acquaintances, and others were associated with vote-switching behaviour in the U.K. General Election of 1992.

16. A sampled individual might misrepresent her preferences in order to manipulate the beliefs of others. I side-step this issue by supposing that information acquisition occurs over an extended period prior to an election, when individuals in the community have little opportunity to hide their true colours. A second issue is that if a voter learns about the signals of others, then she also learns directly about their likely voting behaviour. If the size of the electorate were small, then this effect would be important. Here, however, the electorate is large, and hence knowledge of the likely voting decisions of a few others has only a small impact on the relative likelihood of different pivotal outcomes. Rather than address this issue directly by incorporating the sampling of other voters into the model specification, I side-step by assuming that voters use the sampled preferences of others only to form beliefs about η .

17. In the 1997 U.K. General Election, the 47 nationwide polls contrasted with only 29 constituency-level polls taken in only 26 out of 659 possible constituencies (Evans, Curtice and Norris, 1998).



Notes: Panel (a) relates the size of a private opinion poll m and the accuracy of a voter's beliefs α . Panel (b) uses the community-effect specification to illustrate the false-consensus effect; $\Pr[\hat{y}_i < 0 \mid \bar{u}_i < 0]$ is the probability that a voter who prefers candidate 2 believes him to be more popular, even though he is not.

FIGURE 2
Information, accuracy, and mistaken beliefs

equivalent to specifying $m \approx 42$. Hence, when voters make careful observations of common and accurate information sources, a large value for m would be appropriate.

3.1.3. Community effects. If voters sample individuals who are similar to them, the usefulness of their signals (as measured by m or α) may fall. Suppose that voters learn only from members of their own community, where a community is a small subset of the electorate. Even a large sample will be unable to eliminate any community-specific shock. Formally, suppose that a fraction ω of each idiosyncratic component ε_i is shared with the community:

$$\varepsilon_i = \theta + \tilde{\varepsilon}_i \quad \Rightarrow \quad \tilde{u}_i = \eta + \theta + \tilde{\varepsilon}_i,$$

where $\text{var}[\theta \mid \eta] = \omega\xi^2$ and $\text{var}[\tilde{\varepsilon}_i \mid \eta] = (1 - \omega)\xi^2$. Thus preferences combine a common component across the electorate with a community component and then a further individual component. The parameter ω represents the relative importance of the community effect, and θ arises from community-specific effects. θ cannot be “averaged out” from a community-based sample, and hence beliefs must have variance of at least $\kappa^2 = \omega\xi^2$, equivalent to an electorate-wide sample size of $m = 1/\omega$, and an accuracy of $\alpha = \Phi(\Phi^{-1}(\pi)/\sqrt{\omega})$. If $\omega = 0.2$ so that 20% of individual variation is due to variation across communities, then $m = 5$ and the accuracy for $\pi = 0.6$ would be $\alpha \approx 71\%$.

Lower values for m imply greater correlation between preferences and beliefs. To see this, consider a community-based electorate where $\eta > 0$ but $\tilde{u}_i < 0$. When ω is large, it is highly likely that voter i is drawn from a community where $\eta + \theta < 0$, and so she will believe that her favourite is more popular. An exaggerated confidence (the “false-consensus effect” of Ross, Greene and House (1977) in the psychology literature) in a preferred candidate is in no way irrational (Goeree and Grosser, 2003). If community variation is large then one might expect to see extensive misplaced confidence in preferred candidates (Figure 2(b)).

3.2. Comparative statics

Attention now turns to the determinants of strategic voting.

3.2.1. Strategic incentives. In the unique equilibrium $v^*(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + b^*\hat{\eta} \geq 0]$, $b^*\hat{\eta}$ represents the incentive to vote strategically. Hence the electorate-wide average incentive is

$$\lambda^* \equiv \mathbb{E} \left[\log \frac{f(\gamma \mid \hat{\eta}, v^*)}{f(1-\gamma \mid \hat{\eta}, v^*)} \right] = b^*\eta = b^* \times \zeta \times \Phi^{-1}(\pi). \tag{10}$$

If candidate 1 is strong, so that π is large, voters receive favourable signals of his support, and so the average strategic incentive is large. The determinants of b^* may be ascertained via inspection of $\hat{b}(b)$. Following the social communication story (Section 3.1),

$$\hat{b}(b) = \frac{2\Phi^{-1}(\gamma)\sqrt{m^2 + m(b^2 + 2b)}}{\zeta(1+b)}. \tag{11}$$

If a parameter pushes $\hat{b}(b)$ upward, then (since $\hat{b}(b)$ is decreasing) the fixed point b^* must increase. By inspection, a voter’s response to her beliefs and her average strategic incentive increase with the information available m and with the required-qualified majority γ . Since γ is the “safety” of a disliked status quo, voters rise to a greater need for strategic voting.

The effects of increasing ζ^2 are more complex. When $\eta > 0$, equation (10) reveals that the average signal increases; however, from equation (11), the response to these signals (via b^*) is more sluggish. These two effects work against each other. This is not surprising, since under the social communication specification, the accuracy of beliefs $\alpha = \Phi(\sqrt{m} \times \Phi^{-1}(\pi))$ is invariant to ζ . For a “decision theoretic” voting strategy (see Proposition 2) a (perhaps naive) voter expects sincere actions from others, and the two effects cancel.

Proposition 4. *Consider the naive voting strategy $v^\dagger(\tilde{u}, \hat{\eta}) = \mathbb{I}[b^\dagger\hat{\eta} + \tilde{u} \geq 0]$ from Proposition 2. The average strategic incentive $\lambda^\dagger \equiv b^\dagger\eta$ satisfies $\lambda^\dagger = 2\Phi^{-1}(\gamma)\Phi^{-1}(\pi)m$.*

Turning attention back to the strategic-voting equilibrium, a third effect of idiosyncrasy is present. Increased idiosyncrasy means that, for a fixed strategic-voting incentive, a voter i is less likely to vote strategically, since she is less likely to be relatively indifferent between the candidates. Voter 0 is now less worried about others voting strategically and so is more willing to vote strategically herself. (This is a consequence of the negative-feedback effect.) The net effect of idiosyncrasy is that the average strategic incentive increases with ζ^2 .

Proposition 5. *In equilibrium, the average incentive to vote strategically increases with the qualified majority γ , the underlying popularity of the leading candidate π , the precision of voters’ beliefs m (equivalently, accuracy α), and the idiosyncrasy of the electorate ζ^2 . Also,*

$$\lim_{m \rightarrow \infty} \frac{b^*\eta}{\sqrt{m}} = \Phi^{-1}(\gamma)\Phi^{-1}(\pi)\sqrt{2 + 2\frac{\sqrt{(\Phi^{-1}(\gamma))^2 + \zeta^2}}{\Phi^{-1}(\gamma)}}, \tag{12}$$

so that, asymptotically, the average strategic-voting incentive increases with \sqrt{m} .

Propositions 4 and 5 reveal the rate at which incentives change with more information. When voters are game-theoretic (they anticipate strategic voting by others, rather than assuming that others vote sincerely) they react much more slowly to increased information: as $m \rightarrow \infty$ equilibrium incentives increase (at least asymptotically) with \sqrt{m} rather than m .

3.2.2. Duverger’s law. Allowing $m \rightarrow \infty$ (equivalently, $\alpha \rightarrow 1$ or $\kappa^2 \rightarrow 0$) generates a benchmark case where voters are almost perfectly informed. Hence (for $\pi > 1/2$) almost all voters successfully coordinate on the first candidate, so that the unique strategic-voting equilibrium

is (almost) perfectly Duvergerian when voters have (almost) perfect knowledge of the electoral situation. In contrast, Palfrey (1989) and Myerson and Weber (1993) found multiple Duvergerian equilibria. Heuristically, letting $m \rightarrow \infty$ recovers voters' knowledge of the electoral situation, but removes the assumption that this is common knowledge: dropping the common knowledge that underpins earlier work permits an equilibrium-selection argument.¹⁸

3.2.3. Distance from contention and marginality. The qualified-majority voting model is designed to give insight into behaviour in multi-candidate plurality-rule elections. Under the plurality rule, a common intuition (Cain, 1978, p. 644, for instance) suggests that strategic voting should be greater in marginal elections, and when a preferred candidate is far from contention. The marginality hypothesis stems from the idea that pivotal events are more likely in marginal elections. This is problematic, since such absolute probabilities have no role here; it is the *relative* probability of different pivotal events that matters.

To make progress, I turn to a three-candidate plurality-rule election where the true underlying supports are $\psi_0 + \psi_1 + \psi_2 = 1$. Candidate 0 represents the disliked status quo, such as Buckley in the 1970 New York case. So long as $\frac{1}{3} < \psi_0 < \frac{1}{2}$, this scenario maps back into a qualified-majority voting game: the qualified majority of anti-status-quo voters required to defeat $j = 0$ is $\gamma = \psi_0 / (1 - \psi_0)$; the true underlying support for candidate 1 among them is $\pi = \psi_1 / (1 - \psi_0)$. Using this notation, the ideas of marginality and distance from contention may be formalized. Consider an election in which at least some strategic voting is needed to defeat $j = 0$, so that $\psi_0 > \psi_1 > \psi_2$, or equivalently $\gamma > \pi$. Then I may define

$$\begin{array}{l} \text{Winning} \\ \text{margin} \end{array} \equiv w \equiv \psi_0 - \psi_1 \quad \text{and} \quad \begin{array}{l} \text{Distance from} \\ \text{contention} \end{array} \equiv d \equiv \psi_1 - \psi_2.$$

These parameters are sufficient to determine the electoral situation. In fact,

$$\left. \begin{array}{l} \psi_0 = (1 + 2w + d)/3 \\ \psi_1 = (1 - w + d)/3 \\ \psi_2 = (1 - w - 2d)/3 \end{array} \right\} \Rightarrow \gamma = \frac{1 + 2w + d}{2 - 2w - d} \quad \text{and} \quad \pi = \frac{1 - w + d}{2 - 2w - d}.$$

Simple algebra confirms that both γ and π are increasing in both d and w .

Proposition 6. *When coordination is required to defeat $j = 0$ ($\psi_0 > \psi_1 > \psi_2$) the incentive to vote strategically increases with both the winning margin and the distance from contention.*

This prediction runs against established intuition. After controlling for the distance from contention, strategic voting should be *lower* in more marginal elections.

3.3. The impact of strategic voting

In equilibrium, changes in ζ^2 influence the strategic incentive faced by voters. They also influence voter preferences, however, and hence potentially determine a voter's commitment to her most preferred candidate. To assess the final impact of strategic voting, I must consider both effects.

18. For finite m , the multi-candidate support arising in equilibrium is *not* related in any way to the non-Duvergerian equilibria of common-knowledge voting games. The latter class of equilibria require *exact* knowledge of underlying candidate-support levels and the electorate size, and involve "wrong way" switching.

3.3.1. Impact. Consider the probability that an individual in an appropriate risk population votes strategically. By “risk population” I mean those for whom a strategic vote makes sense. When $\pi > 1/2$, a voter who prefers candidate 2 (so that $\tilde{u}_i < 0$) but (correctly) believes that candidate 1 is best placed to win (so that $\hat{\eta}_i > 0$) is “at risk” of voting strategically.

An easy way to study an at-risk individual is to equip her with beliefs $\hat{\eta}_i = \eta$ (so that she has a correct expectation of the electoral situation; of course, she does not know that this is the case) and then examine her behaviour. She votes for candidate 1 with probability $\hat{p} = \Pr[(1+b)\eta + \varepsilon_i \geq 0 | \eta] = \Phi((1+b)\Phi^{-1}(\pi))$.¹⁹ Of course, this voter actually prefers candidate 1 with probability π , and so a strategic vote is observed with probability $\hat{p} - \pi$. Furthermore, she is at risk of voting strategically when she prefers candidate 2 (with probability $1 - \pi$). Hence, the impact of strategic voting, measured as a proportion of the risk population, is

$$\frac{\hat{p} - \pi}{1 - \pi} = \frac{\Phi((1+b)\Phi^{-1}(\pi)) - \pi}{1 - \pi}. \quad (13)$$

Inspection of this expression, coupled with earlier results, yields the following proposition.

Proposition 7. *In the unique strategic-voting equilibrium, the probability that an at-risk individual votes strategically increases with the required qualified majority, the strength of the leading candidate, and the information available to voters, but decreases with idiosyncrasy.*

3.3.2. Political impact. When will strategic voting change the election outcome? Since a voter prefers candidate 1 with probability π , but votes for him with probability p , strategic voting changes the election result when $p > \gamma > \pi$. Now, straightforward manipulations yield

$$\pi = \Phi\left(\frac{\Phi^{-1}(p)}{1+b} \sqrt{\frac{(m-1) + (1+b)^2}{m}}\right). \quad (14)$$

Given m , ζ^2 , and m , equation (14) can be used to remove the effect of strategic voting, and “invert” an election. Equations (13) and (14) are exploited in the next section.

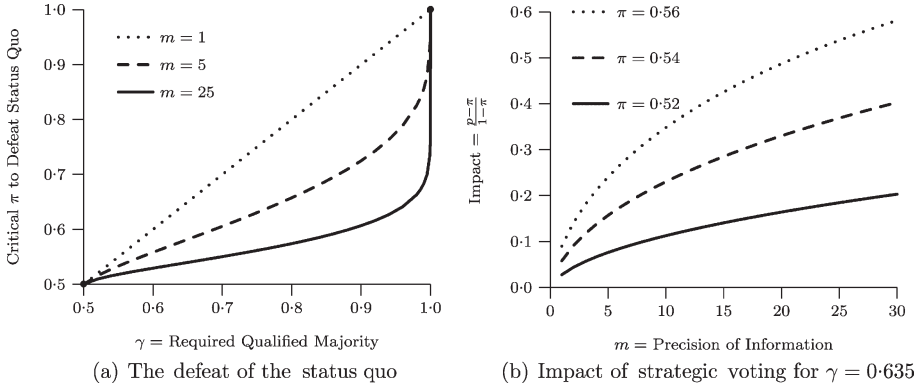
4. APPLICATIONS

The qualified-majority voting model does not provide a full description of all multi-candidate plurality elections, but nevertheless gives insights into scenarios in which supporters of two candidates wish to coordinate against a disliked third candidate. In “beat the conservative” settings, the model is rich enough to allow “back of the envelope” calibration exercises.

4.1. The (political) impact of strategic voting

To use equations (13) and (14), a value for ζ^2 is needed. This may be pinned down by considering the preference intensity of a candidate’s core supporters: if the median supporter of candidate 1 likes candidate 1 twice as much as candidate 2, then $\Pr[\tilde{u}_i \geq \log 2] = \pi/2$. Together with π , this

19. Here I have actually given her an artificial belief $\hat{\eta}_i = \eta$, and so $\varepsilon_i \sim N(0, \zeta^2)$. If I had considered her behaviour conditional on her actually holding a belief $\hat{\eta}_i$, then the distribution of ε_i changes slightly. The reason is that \tilde{u}_i and hence ε_i is a component of $\hat{\eta}_i$, and so knowledge of $\hat{\eta}_i$ influences the (conditional) distribution of ε_i . Correcting for this has a minimal effect on the analysis (Appendix B.4).



Notes: Panel (a) illustrates the political impact of strategic voting. For each m , combinations of π and γ above the line result in the defeat of the status quo. Panel (b) relates the precision of information to the impact of strategic voting for values of γ and π that are consistent with the 1970 New York Senatorial Election.

FIGURE 3
The (political) impact of strategic voting

determines ζ^2 . In fact, for $\pi = 1/2$ it solves for $\zeta^2 = 1.056$ (see Appendix B.5), and here I fix ζ^2 at this value.

With ζ^2 fixed, the political impact of strategic voting can be calculated. Figure 3(a) illustrates combinations of γ and $\pi > \frac{1}{2}$ which ensure the defeat of the status quo in equilibrium. For instance, when $\pi > \gamma$ then the first candidate wins even when there is no strategic voting. Hence in the region above the 45° line the status quo is always defeated. Turning to the region below the 45°, consider the hatched line for the case of $m = 5$. For (γ, π) combinations below the hatched line, $\gamma > p > \pi$, and hence the status quo prevails. On the other hand, for (γ, π) combinations above the hatched line, $p > \gamma > \pi$, and hence the status quo is defeated even though it would prevail if everyone voted sincerely.


For a fixed value of γ , I may calculate the direct impact of strategic voting (*i.e.* the probability that an “at risk” individual votes strategically) for various values of π and m . Figure 3(b) displays the results of such an exercise for $\gamma = 0.635$; observe how the impact of strategic voting rises with the information available to voters.

As a practical illustration, I now turn back to the 1970 New York senatorial election (Table 1). Mapping this example into the model yields the parameters γ and p ,²⁰

$$\gamma = \frac{\psi_0}{\psi_1 + \psi_2} = \frac{2,288,190}{3,605,704} = 0.635, \quad \text{and} \quad p = \frac{\psi_1}{\psi_1 + \psi_2} = \frac{2,171,232}{3,605,704} = 0.602. \quad (15)$$

Hence liberals needed to achieve a qualified majority of $\gamma = 0.635$ to defeat the disliked Buckley and yet achieved only $p = 0.602$; this split in the vote allowed Buckley to win. This discussion suggests that strategic voting might not have changed the election outcome. Nevertheless, equation (14) may be used to calculate the value for π that is consistent with the data. Doing so with $m = 25$, for instance, yields $\pi = 0.528$. This suggests that, absent strategic voting, Ottinger’s true support may have been closer to $0.528 \times 61\% \approx 32.2\%$.

20. Equation (15) involves a slight abuse of notation. In Section 3.2, the parameters ψ_0 , ψ_1 , and ψ_2 were related to the true underlying support parameter π , whereas equation (15) uses the voting probability p .



Notes: This barycentric plot illustrates vote shares for 527 English constituencies. A constituency “•” is a weighted average of the three extremes, with weights corresponding to the major parties’ relative vote shares. The sides of the simplex (solid lines) are where only two parties receive votes, and hence would correspond to “Duvergerian” equilibria. The hatched lines are where two parties tie for the lead, and hence separate the “win zones” for each party. The dotted lines are where two parties tie for second place, and hence would correspond to “non-Duvergerian” equilibria. The two horizontal lines bracket the 270 constituencies in which the Conservative party polled between one-third and half of the votes for major parties, so that anti-Conservative coordination was required to avoid a Tory win.

FIGURE 4
U.K. General Election 1997

4.2. United Kingdom (1997)

The equilibria of a “common knowledge” plurality-voting games (Palfrey, 1989; Myerson and Weber, 1993) predict either the complete coordination of strategic voting, or an exact tie for second place. Clearly, the 1970 New York election deviates from these predictions. The U.K. General Election of 1997 provides a further illustration. In 1997, the three major political parties (Conservative, Labour, and Liberal Democrat) competed throughout 527 English parliamentary constituencies. Vote shares are plotted in Figure 4; by inspection, Duvergerian and non-Duvergerian equilibrium outcomes are absent.

The strategic-voting equilibria described in this paper are consistent with multi-candidate support. As an exercise, the model can be calibrated against the U.K. election outcomes. Fixing ξ^2 as above, a value for m (equivalently, α or ω) is required. The “community effects” discussed in Section 3.1 suggest that, when relying on social communication as an information source, the accuracy of voters’ beliefs may be limited by cross-community shocks. To take an example, suppose that $\omega = 1/5$, or equivalently $m = 5$. If $\pi = 0.635$, then a voter correctly identifies the leading challenger with probability 78%, and hence 22% of the electorate are mistaken in their perception of the leading challenger.

Does this resonate with observation? The 1997 British Election Survey reveals that 68.5% of voters who expected their preferred party to come second actually found that it came third (Fisher, 2000); an instance of the false-consensus effect. Similarly, roughly half of those whose preferred party came third expected it to come second. (Brown (1982) and Baker, Koestner, Worren, Losier and Vallerand (1995) reported similar findings in U.S. and Canadian elections.) The British

TABLE 2
Inverted election results: English Constituencies 1997

	Actual	$\omega = 0.05$	$\omega = 0.1$	$\omega = 0.2$	$\omega = 0.3$	$\omega = 0.4$	$\omega = 0.5$
Lab > Con > Lib	118	57	70	82	105	105	105
Con > Lab > Lib	73	134	121	109	86	86	86
Con > Lib > Lab	54	70	68	65	65	63	59
Lib > Con > Lab	25	9	11	14	14	16	20
Con \rightarrow Lab	—	61	48	36	13	13	13
Con \rightarrow Lib	—	16	14	11	11	9	5
Con seats lost	—	77	62	47	24	22	18
Impact (%)	—	35.8	31.9	26.3	21.9	18.1	14.6

Notes: The first four rows classify the 270 seats according to the relative vote shares of the three main parties. I have excluded constituencies in which the Conservative party came third, hence there are four possible configurations. The first column gives the classification according to the actual results. The remaining columns give the results when the strategic-voting element is removed, for values of ω from 0.05 to 0.5 (and hence m from 2 to 20). The fifth, sixth, and seventh rows give the number of seats lost by the Conservatives due to strategic switching. The final row gives the probability of a strategic vote by a typical member of the risk population; that is, someone who prefers the trailing challenger and is equipped with a correct signal realization indicating the electoral situation (equation (13)). Con, Conservative; Lab, Labour; Lib, Liberal.

voters' poor understanding of the electoral situation coincides with the community-effect model (see Figure 2(b)). Thus a situation in which voters' beliefs are somewhat "fuzzy" generates (at first blush) the features observed in survey data.

Returning to the calibration exercise, I consider the 270 constituencies where the Conservatives polled between $\frac{1}{3}$ and $\frac{1}{2}$ of the vote, and identify them as the status quo.²¹ For each constituency, I calculated appropriate values for γ and p . For a variety of values of ω (and hence m) I calculated b^* and, using equation (14), "inverted" the result to obtain the notional true level of support for each challenging candidate. Using these notional levels of support, I then calculated the election results in the absence of strategic voting.²²

This exercise should be seen as nothing more than indicative. Nevertheless, the results are interesting (Table 2). For $\omega = 0.2$, so that community effects account for 20% of preference idiosyncrasy, the results suggest that the Conservative party lost 47 seats due to strategic voting. This same parameter choice suggests that, according to the theory, strategic voting affected around 26% of the "risk population" of trailing candidate supporters. This matches the results of Fisher (2000). He estimated (using British Election Survey (BES) data, and all English constituencies) that 24.4% of votes in the risk population of third-party supporters voted strategically in 1997. The incidence of strategic voting from a calibrated model appears to be about right. Moving beyond calibration exercises, the model's comparative-static predictions that may be taken to the data. Proposition 6 predicts that the incentive to vote strategically should increase with both the distance from contention and the winning margin. Empirically, the distance from contention is well known as a strong predictor of strategic voting (Niemi, Whitten and Franklin, 1992, 1993; Evans and Heath, 1993; Franklin, Niemi and Whitten, 1994; Heath and Evans, 1994). The predicted effect of the winning margin, however, clashes with the intuition of Cain (1978) and others. In recent work, Fisher (2000) found that strategic voting increases with the winning margin in U.K. General Elections. The effect is weak, but significant at the 10% level

21. The implicit assumption is that supporters of the Labour and Liberal Democrat parties rank the Conservatives last. In fact a significant proportion of Liberal Democrat voters ranked the Labour party last in 1997. Nevertheless, this simplification is used to generate "ball park" indications.

22. Myatt and Fisher (2002a) conducted a similar exercise using a model with two different preference types.

when estimated across the three elections of 1987, 1992, and 1997.²³ This means that there is empirical support for the present theory relative to informal intuition.

5. CONCLUDING REMARKS

Duverger (1954) claimed that “simple-majority single-ballot system favours the two-party system”, his “psychological effect” leading voters to abandon a trailing candidate, so that (Droop, 1871) “an election is usually reduced to a contest between the two most popular candidates ... votes will be thrown away, unless given in favour of one or other of the parties between whom the election really lies”. The models of Palfrey (1989), Myerson and Weber (1993), and Cox (1994) support a strictly Duvergerian claim. They rely, however, on common knowledge of the candidates “between whom the election really lies”.

In this paper I have removed the common-knowledge assumption. The unique strategic-voting equilibrium yields multi-candidate support, with only partial coordination of the electorate. Whereas they are very different from the non-Duvergerian equilibria that were important to Cox (1994, 1997), strategic-voting equilibria precisely incorporate the intuition that he offered. For instance, Cox (1997, p. 86) interpreted the result for the Ross and Cromarty constituency (where the Liberal and Labour candidates almost tied, permitting a Conservative win) from the 1970 U.K. General Election as a potential non-Duvergerian equilibrium. He claimed that this interpretation required that “it was not clear who was in third and who in second” and that “neither [challenger] suffered from strategic desertion”.

A rich set of comparative statics related strategic voting to the underlying electoral situation, preference characteristics, and the accuracy of voters’ information sources. Some of these may seem surprising: the negative-feedback effect suggests that voters should be more cautious in their behaviour when they expect others to act strategically, and the comparative-static results reject the marginality hypothesis. While I have attempted no real empirical analysis, a calibration of the model applied to the U.K. General Election of 1997 resonates with observed behaviour and offers insights into the wider consequences of strategic voting. Moreover, related work takes some of the predictions to the data, and they are not rejected.

APPENDIX A. OMITTED PROOFS

Proof of Lemma 1. Follows from Bayesian updating (DeGroot, 1970). ||

Proof of Lemma 2. Multi-candidate support and monotonicity ensure that $v(\tilde{u}_i, \hat{\eta}_i) = 1$ for $(\tilde{u}_i, \hat{\eta}_i)$ large enough, and so $\lim_{\eta \uparrow \infty} E[v(\tilde{u}_i, \hat{\eta}_i) | \eta] = 1$. Similarly, $\lim_{\eta \downarrow -\infty} E[v(\tilde{u}_i, \hat{\eta}_i) | \eta] = 0$. The two conditions also ensure that $E[v(\tilde{u}_i, \hat{\eta}_i) | \eta]$ is strictly increasing in η . A voter’s posterior over η yields a positive density on \mathbb{R} , and so $f(p | \hat{\eta}, v)$ has full support. Continuity follows from the properties of the expectation with respect to a continuous density. ||

Proof of Lemma 3. $H(\eta) \equiv E[v(\hat{\eta}_i, \tilde{u}_i) | \eta]$ is smoothly increasing in η . Furthermore,

$$F(p | \hat{\eta}, v) = \Pr[\eta \leq H^{-1}(p) | \hat{\eta}] = \Phi\left(\frac{H^{-1}(p) - \hat{\eta}}{\kappa}\right), \quad (16)$$

following from the normality of posterior beliefs over η . Write $h(\eta) = H'(\eta)$ and differentiate:

$$f(p | \hat{\eta}, v) = \frac{1}{h(H^{-1}(p))\kappa} \phi\left(\frac{H^{-1}(p) - \hat{\eta}}{\kappa}\right) \quad (17)$$

23. Myatt and Fisher (2002b) studied a decision-theoretic three-party model in which voters’ beliefs take a Dirichlet form over ψ_0 , ψ_1 , and ψ_2 and where the comparative-static predictions coincide with the present paper. An appropriate strategic-incentive variable explains the pattern of strategic voting across all three elections. Fisher (2001) included a range of other explanatory variables, including political interest, education, party identification, and local campaign-spending and obtained similar results.

where $\phi(\cdot)$ is the density of the standard normal. Applying this,

$$\begin{aligned} \log \frac{f(\gamma \mid \hat{\eta}, v)}{f(1-\gamma \mid \hat{\eta}, v)} &= \log \frac{h(H^{-1}(1-\gamma))}{h(H^{-1}(\gamma))} - \frac{(H^{-1}(\gamma) - \hat{\eta})^2}{2\kappa^2} + \frac{(H^{-1}(1-\gamma) - \hat{\eta})^2}{2\kappa^2} \\ &= \log \underbrace{\frac{h(H^{-1}(1-\gamma))}{h(H^{-1}(\gamma))} + \frac{(H^{-1}(1-\gamma))^2 - (H^{-1}(\gamma))^2}{2\kappa^2}}_{\text{Intercept}=a} + \underbrace{\frac{H^{-1}(\gamma) - H^{-1}(1-\gamma)}{\kappa^2}}_{\text{Slope}=b} \times \hat{\eta}. \end{aligned} \tag{18}$$

By inspection, this is linear in the voter’s belief type $\hat{\eta}$.

Now suppose that voters $i \in \{1, \dots, n\}$ use a linear strategy $v(\tilde{u}_i, \hat{\eta}_i) = \mathbb{I}[\tilde{u}_i + a + b\hat{\eta}_i \geq 0]$. Hence, $v_i = 1$ if and only if $a + (1 + b)\eta \geq -[(\tilde{u}_i - \eta) + b(\hat{\eta}_i - \eta)]$. Conditional on η , the R.H.S. of this inequality is normally distributed with zero expectation and variance $\tilde{\kappa}^2 \equiv \text{var}[(\tilde{u}_i - \eta) + b(\hat{\eta}_i - \eta)]$. Thus the probability p of a vote for candidate 1 satisfies

$$p = H(\eta) = \Phi\left(\frac{a + (1 + b)\eta}{\tilde{\kappa}}\right) \Rightarrow \eta = H^{-1}(p) = \frac{\tilde{\kappa}\Phi^{-1}(p) - a}{1 + b}, \tag{19}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal. Differentiating,

$$h(\eta) = H'(\eta) = \frac{1 + b}{\tilde{\kappa}} \phi\left(\frac{a + (1 + b)\eta}{\tilde{\kappa}}\right) \Rightarrow h(H^{-1}(p)) = \frac{1 + b}{\tilde{\kappa}} \phi\left(\Phi^{-1}(p)\right). \tag{20}$$

Begin with the first term of equation (18). Firstly employ the symmetry of the normal distribution to note that $\Phi^{-1}(1 - \gamma) = -\Phi^{-1}(\gamma)$, and that $\phi(z) = \phi(-z)$. It follows that

$$\log \frac{h(H^{-1}(1-\gamma))}{h(H^{-1}(\gamma))} = \log \frac{\phi(\Phi^{-1}(1-\gamma))}{\phi(\Phi^{-1}(\gamma))} = 0. \tag{21}$$

Next consider the second term of equation (18),

$$(H^{-1}(\gamma))^2 = \frac{(\tilde{\kappa}\Phi^{-1}(\gamma) - a)^2}{(1 + b)^2} = \frac{[\tilde{\kappa}\Phi^{-1}(\gamma)]^2 + a^2 - 2a\tilde{\kappa}\Phi^{-1}(\gamma)}{(1 + b)^2}. \tag{22}$$

Similarly,

$$(H^{-1}(1-\gamma))^2 = \frac{[\tilde{\kappa}\Phi^{-1}(1-\gamma)]^2 + a^2 - 2a\tilde{\kappa}\Phi^{-1}(1-\gamma)}{(1 + b)^2} = \frac{[\tilde{\kappa}\Phi^{-1}(\gamma)]^2 + a^2 + 2a\tilde{\kappa}\Phi^{-1}(\gamma)}{(1 + b)^2}, \tag{23}$$

and so it follows that

$$\frac{H^{-1}(1-\gamma)^2 - H^{-1}(\gamma)^2}{2\kappa^2} = \frac{2a\tilde{\kappa}\Phi^{-1}(\gamma)}{\kappa^2(1 + b)^2}. \tag{24}$$

The final term is simply

$$\frac{H^{-1}(\gamma) - H^{-1}(1-\gamma)}{\kappa^2} \hat{\eta} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)(1 + b)}{(1 + b)^2\kappa^2} \hat{\eta}. \tag{25}$$

Assemble these terms and obtain

$$\begin{aligned} \log \left[\frac{f(\gamma \mid \hat{\eta}, v)}{f(1-\gamma \mid \hat{\eta}, v)} \right] &= \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)(a + (1 + b)\hat{\eta})}{\kappa^2(1 + b)^2} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)a}{\kappa^2(1 + b)^2} + \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)\hat{\eta}}{\kappa^2(1 + b)} \\ &= \frac{a\hat{b}}{(1 + b)} + \hat{b}\hat{\eta}, \end{aligned} \tag{26}$$

where $\hat{b} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)}{\kappa^2(1+b)}$. The slope parameter depends upon $\tilde{\kappa}$. Recalling that $\text{cov}[\tilde{u}_i, \hat{\eta}_i \mid \eta] = \kappa^2$ (see Lemma 1), $\tilde{\kappa}^2 = \text{var}[(\tilde{u}_i - \eta) + b(\hat{\eta}_i - \eta) \mid \eta] = \xi^2 + b^2\kappa^2 + 2b\kappa^2$, and hence

$$\hat{b}(b) = \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + (b^2 + 2b)\kappa^2}}{\kappa^2(1+b)}. \quad \parallel$$

The following lemma is used in the proofs of Propositions 1 and 5.

Lemma 4. *The mapping $\hat{b}(b)$ has a unique fixed point $b^* > 0$. This is bounded as follows*

$$2\Phi^{-1}(\gamma) \leq \kappa b^* \leq \Phi^{-1}(\gamma) \sqrt{2 + 2\sqrt{\frac{(\Phi^{-1}(\gamma))^2 + \xi^2}{(\Phi^{-1}(\gamma))^2}}}, \quad (27)$$

where κb^* attains the upper bound as $\kappa^2 \rightarrow 0$.

Proof. The uniqueness of b^* follows from the fact that $\hat{b}(b)$ is decreasing:

$$\hat{b}'(b) = \frac{2\Phi^{-1}(\gamma)}{\kappa^2} \left\{ \frac{(1+b)\kappa^2}{(1+b)\sqrt{\xi^2 + (b^2 + 2b)\kappa^2}} - \frac{\sqrt{\xi^2 + (b^2 + 2b)\kappa^2}}{(1+b)^2} \right\} \quad (28)$$

$$= \hat{b}(b) \left\{ \frac{(1+b)\kappa^2}{\xi^2 + (b^2 + 2b)\kappa^2} - \frac{1}{1+b} \right\}. \quad (29)$$

For $b \geq 0$, this derivative is negative if and only if

$$1 \geq \frac{(1+b)^2\kappa^2}{\xi^2 + (b^2 + 2b)\kappa^2} = \frac{(b^2 + 2b + 1)\kappa^2}{\xi^2 + (b^2 + 2b)\kappa^2} \Leftrightarrow \kappa^2 \leq \xi^2, \quad (30)$$

which holds since ξ^2 is an upper bound to the variance of a voter's signal. The lower bound to b^* is simply $\lim_{b \rightarrow \infty} \hat{b}(b)$. To obtain an upper bound for the fixed point, write $\hat{b}(b)$ as

$$\hat{b}(b)\kappa = \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + (b^2 + 2b)\kappa^2}}{\kappa + b\kappa}. \quad (31)$$

Make the change of variable $\beta = \kappa b$ to obtain

$$\hat{\beta}(\beta) = \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + \beta^2 + 2\kappa\beta}}{\kappa + \beta} = 2\Phi^{-1}(\gamma)\sqrt{\frac{\xi^2 + \beta^2 + 2\kappa\beta}{\kappa^2 + \beta^2 + 2\kappa\beta}}. \quad (32)$$

It is clear that the R.H.S. is decreasing in κ . Hence sending $\kappa \downarrow 0$,

$$\hat{\beta}(\beta) \leq \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + \beta^2}}{\beta}. \quad (33)$$

To obtain the desired upper bound from this, solve the equation $\beta^4 - (2\Phi^{-1}(\gamma))^2(\xi^2 + \beta^2) = 0$. This is quadratic in β^2 , and may be solved to obtain the positive root

$$\beta^2 = \frac{(2\Phi^{-1}(\gamma))^2}{2} \left\{ 1 + \sqrt{1 + \frac{\xi^2}{(\Phi^{-1}(\gamma))^2}} \right\}. \quad (34)$$

Taking the square root yields the upper bound; clearly, this is attained as $\kappa \downarrow 0$. \parallel

Proof of Proposition 1. From Lemma 3, an equilibrium strategy must be linear, and satisfy $b^* = \hat{b}(b^*)$ and $a^* = \hat{a}(a^*, b^*)$. From Lemma 4, $\hat{b}(b)$ has a unique positive fixed point b^* . Inspecting $\hat{a}(a, b^*)$, the equilibrium intercept must be $a^* = 0$. Comparative statics follow since, by inspection, the mapping $\hat{b}(b)$ is increasing in γ but decreasing in κ^2 and ξ^2 . \parallel

Proof of Proposition 2. Evaluate $\hat{b}(b)$ and $\hat{a}(a, b)$ for $a = b = 0$. \parallel

Proof of Proposition 3. Incorporating the public signal μ , equation (18) becomes

$$\begin{aligned} \log \frac{f(\gamma \mid \hat{\eta}, \mu, v)}{f(1-\gamma \mid \hat{\eta}, \mu, v)} &= \log \frac{h(H^{-1}(1-\gamma))}{h(H^{-1}(\gamma))} + \frac{(H^{-1}(1-\gamma))^2 - (H^{-1}(\gamma))^2}{2\text{var}[\eta \mid \hat{\eta}, \mu]} \\ &\quad + \frac{H^{-1}(\gamma) - H^{-1}(1-\gamma)}{\text{var}[\eta \mid \hat{\eta}, \mu]} \times E[\eta \mid \hat{\eta}, \mu]. \end{aligned} \tag{35}$$

$E[\eta \mid \hat{\eta}, \mu]$ is linear in $\hat{\eta}$, and hence so is log-likelihood ratio. In fact, using the “ ζ ” notation,

$$E[\eta \mid \hat{\eta}, \mu] = \frac{\mu + \zeta \hat{\eta}}{1 + \zeta}, \quad \text{and} \quad \text{var}[\eta \mid \hat{\eta}, \mu] = \frac{\zeta \kappa^2}{1 + \zeta}. \tag{36}$$

$H(\eta)$ takes the form used in the proof of Lemma 3, and so substituting yields

$$\begin{aligned} \log \frac{f(\gamma \mid \hat{\eta}, \mu, v)}{f(1-\gamma \mid \hat{\eta}, \mu, v)} &= \frac{2a\bar{\kappa}\Phi^{-1}(\gamma)(1+\zeta)}{\zeta\kappa^2(1+b)^2} + \frac{2\bar{\kappa}\Phi^{-1}(\gamma)(\mu + \zeta\hat{\eta})}{\zeta\kappa^2(1+b)} \\ &= \frac{2a\bar{\kappa}\Phi^{-1}(\gamma)(1+\zeta)}{\zeta\kappa^2(1+b)^2} + \frac{2\bar{\kappa}\Phi^{-1}(\gamma)\mu}{\zeta\kappa^2(1+b)} + \frac{2\bar{\kappa}\Phi^{-1}(\gamma)\hat{\eta}}{\kappa^2(1+b)} = \frac{\hat{b}(1+\zeta)a}{\zeta(1+b)} + \frac{\hat{b}\mu}{\zeta} + \hat{b}\hat{\eta}. \end{aligned} \tag{37}$$

Setting $b = \hat{b}(b) = b^*$ yields equation (7), which solves to yield a^* . \parallel

Proof of Proposition 4. Calculate $b^\dagger = \hat{b}(0)$ using equation (11). \parallel

Proof of Proposition 5. The effects of γ , π , and m follow by inspection. For ζ^2 ,

$$b^* = \hat{b}(b^*) \Rightarrow \frac{db^*}{d\zeta} = \frac{\partial \hat{b}(b^*)}{\partial \zeta} + \hat{b}'(b^*) \frac{db^*}{d\zeta} \Rightarrow \frac{db^*}{d\zeta} = \frac{1}{1 - \hat{b}'(b^*)} \frac{\partial \hat{b}(b^*)}{\partial \zeta}.$$

Equation (11) yields $\partial \hat{b}(b^*) / \partial \zeta = -\hat{b}(b^*) / \zeta = -b^* / \zeta$. Combine the two expressions to obtain

$$\frac{d[b^*\zeta]}{d\zeta} = b^* + \zeta \frac{db^*}{d\zeta} = b^* \left[1 - \frac{1}{1 - \hat{b}'(b^*)} \right] = -\frac{b^*\hat{b}'(b^*)}{1 - \hat{b}'(b^*)} > 0.$$

The last inequality follows from $\hat{b}'(b^*) < 0$, since $\hat{b}(b^*)$ is decreasing, and $\hat{b}'(b^*) > -1$, since b^* is a stable fixed point (Appendix B.3). It remains to consider the behaviour of strategic incentives as $m \rightarrow \infty$. The limiting behaviour of the equilibrium strategic incentive as $m \rightarrow \infty$ is equivalent to that when $\kappa^2 \rightarrow 0$. The proof to Lemma 4 demonstrates that κb^* attains the upper bound of equation (27) as $\kappa \rightarrow 0$. This completes the proof. \parallel

Proof of Proposition 6. By inspection γ is increasing in both w and d and π is increasing in d . Differentiation of π with respect to w demonstrates that π is also increasing in w . \parallel

Proof of Proposition 7. By inspection, the impact measure from equation (13) is increasing in π and b^* . In turn, b^* is increasing in γ and m but decreasing in ζ^2 . \parallel

APPENDIX B. OMITTED ANALYSIS

B.1. Generalizing the solution concept and preferences

A strategic-voting equilibrium is a strategy profile that works in large electorates. Of course, there are two symmetric and monotonic strategy profiles that always yield Nash equilibria: the Duvergerian equilibria in which voters ignore their type realizations and all vote for the same candidate. They can form strategic-voting equilibria if the multi-candidate support requirement is dropped. For the following modified definition, $\text{BR}(\bar{u}, \hat{\eta} \mid v)$ is the set of best responses.

Definition 3 (Modified to allow full coordination). A *strategic-voting equilibrium* is a symmetric and monotonic strategy $v^*(\bar{u}, \hat{\eta})$ such that $\Pr[v^*(\bar{u}, \hat{\eta}) \in \lim_{n \rightarrow \infty} \text{BR}(\bar{u}, \hat{\eta} \mid v^*)] = 1$.

The (type independent) strategies $v(\tilde{u}, \hat{\eta}) \equiv 1$ and $v(\tilde{u}, \hat{\eta}) \equiv 0$ fit this revised definition. Nevertheless, Duvergerian equilibria disappear under a more general pay-off specification. To see why, assume that a type realization $(\tilde{u}_i, \hat{\eta}_i)$ leads to pay-offs $u_{i1} = U(\tilde{u}_i)$ and $u_{i2} = 1 - u_{i1}$, where $U(\cdot) : \mathbb{R} \mapsto [0, 1]$ is increasing and continuous. For simplicity, $U(\cdot)$ is symmetric around zero, so that $U(-\tilde{u}_i) = 1 - U(\tilde{u}_i)$. Fixing some $\bar{u} > 0$, assume that $U(\tilde{u}_i) = 0$ for $\tilde{u}_i \leq -\bar{u}$, $U(\tilde{u}_i) = 1$ for $\tilde{u}_i \geq \bar{u}$, and that $U(\tilde{u}_i)$ is strictly increasing and differentiable for $\tilde{u}_i \in (-\bar{u}, \bar{u})$. The specification used in the main paper is obtained by setting $U(\tilde{u}_i) = \frac{e^{\tilde{u}_i}}{1+e^{\tilde{u}_i}}$ and $\bar{u} = \infty$. For $\bar{u} < \infty$, however, the situation is subtly different. When $|\tilde{u}_i| \geq \bar{u}$ it is a weakly dominant strategy to vote sincerely. An elimination of weakly dominated strategies ensures that $v(\tilde{u}_i, \hat{\eta}_i) = \mathbb{I}[\tilde{u}_i \geq 0]$ for $|\tilde{u}_i| \geq \bar{u}$, and so voting strategies automatically exhibit multi-candidate support. Lemma 2 and the first part of Lemma 3 apply. For $|\tilde{u}_i| < \bar{u}$,

$$\text{BR}(\bar{u}, \hat{\eta} \mid v) \rightarrow \mathbb{I} \left[\log \left[\frac{U(\bar{u})}{1 - U(\bar{u})} \right] + \lambda(\hat{\eta}) \geq 0 \right] \quad \text{where} \quad \lambda(\hat{\eta}) = a + b\hat{\eta}.$$

Hence Duvergerian equilibria are eliminated, and strategic-voting equilibria involve linear strategies. In fact, there exists an equilibrium with $a^* = 0$ and $b^* > 0$ where b^* satisfies

$$b^* = \frac{2H^{-1}(\gamma)}{\kappa^2} \quad \text{where} \quad H(\eta) = \int_{-\infty}^{\infty} \Phi \left(\frac{\hat{\eta}_i - U^{-1}([1 + e^{b^*} \hat{\eta}_i]^{-1})}{\sqrt{\xi^2 - \kappa^2}} \right) d\Phi \left(\frac{\hat{\eta}_i - \eta}{\kappa} \right).$$

B.2. Equilibria in finite electorates

The strategic-voting equilibrium concept does not lead to a Nash equilibrium for a fixed n . Adopting the generalized preferences of Section B.1 with $\bar{u} < \infty$, it is possible to construct a symmetric Nash equilibrium with multi-candidate support in a finite electorate. Here I sketch a proof of this claim.

First, some preliminaries. For the generalized-preference model, there is no loss in setting $\xi^2 = 1$. Rather than use $\eta \in \mathbb{R}$, I write $\pi \in [0, 1]$ for the underlying state. For $\pi \in (0, 1)$, and hence $\eta = \Phi^{-1}(\pi)$, types satisfy $\tilde{u}_i \sim N(\eta, 1)$ and $\hat{\eta}_i \sim N(\eta, \frac{1}{m})$. For $\pi \in \{0, 1\}$, voters are all extreme preference types; for instance, when $\pi = 1$ all voters vote for the first candidate.

Recall that $\pi \in [0, 1]$ is the probability that a voter’s true favourite is the first candidate, whereas $p \in [0, 1]$ is the probability that she actually votes for the first candidate. Fixing a voting strategy, these are related via $p(\pi) = E[v(\tilde{u}_i, \hat{\eta}_i) \mid \pi]$. Due to the possibility of extreme preference types, $p(\pi) \geq \Pr[\tilde{u}_i \geq \bar{u} \mid \pi]$ and $p(\pi) \leq \Pr[\tilde{u}_i \geq -\bar{u} \mid \pi]$, and hence

$$1 \geq \Phi(\Phi^{-1}(\pi) + \bar{u}) \geq p(\pi) \geq \Phi(\Phi^{-1}(\pi) - \bar{u}) \geq 0. \tag{38}$$

$p(\pi)$ satisfies $p(0) = 0$ and $p(1) = 1$, and the properties of the expectation ensure that it is smooth. Thus, I focus on the set of smooth functions $P = \{p(\pi) : [0, 1] \mapsto [0, 1]\}$ that arise from $p(\pi) = E[v(\tilde{u}_i, \hat{\eta}_i) \mid \pi]$ for some voting strategy $v(\tilde{u}_i, \hat{\eta}_i)$, and hence satisfy (38). Employing the sup norm, the set of real-valued continuous functions on $[0, 1]$ is a Banach space. P is equicontinuous and bounded, and its closure \bar{P} forms a compact and convex subset of a Banach space.²⁴ I now build a continuous mapping back into \bar{P} . Taking $p(\pi)$, I form the cumulative distribution function

24. P is equicontinuous if, given $\varepsilon > 0$, there is a $\delta > 0$ such that $|\pi_H - \pi_L| \leq \delta \Rightarrow |p(\pi_H) - p(\pi_L)| \leq \varepsilon$ for all $p(\pi) \in P$. To show this, begin by fixing $\varepsilon > 0$ and set $\pi^\dagger = \frac{\Phi(\Phi^{-1}(\varepsilon) - \bar{u})}{2}$. This choice for π^\dagger ensures that $\pi \leq 2\pi^\dagger \Rightarrow p(\pi) \leq \varepsilon$ and $\pi \geq 1 - 2\pi^\dagger \Rightarrow p \geq 1 - \varepsilon$, following equation (38). (As I confirm just below, this ensures that $p(\pi)$ cannot change too much close to the boundaries of the unit interval.) I claim that $p'(\pi)$ is uniformly bounded for all $p(\pi) \in P$ and $\pi \in [\pi^\dagger, 1 - \pi^\dagger]$. Recalling that $\pi = \Phi(\eta)$,

$$p'(\pi) = \frac{d\eta}{d\pi} \times \frac{\partial E[v(\tilde{u}_i, \hat{\eta}_i) \mid \eta]}{\partial \eta} = \frac{1}{\phi(\Phi^{-1}(\pi))} \frac{\partial}{\partial \eta} \int_{\mathbb{R}^2} v(\tilde{u}_i, \hat{\eta}_i) f(\tilde{u}_i, \hat{\eta}_i \mid \eta) d(\tilde{u}_i, \hat{\eta}_i),$$

where $f(\tilde{u}_i, \hat{\eta}_i \mid \eta)$ is a joint normal density. This density is increasing in η if and only if $\eta \leq \hat{\eta}_i$. Hence

$$p'(\pi) \leq \frac{1}{\phi(\Phi^{-1}(\pi))} \int_{\mathbb{R}^2} \mathbb{I}[\hat{\eta}_i \leq \eta] \frac{\partial f(\tilde{u}_i, \hat{\eta}_i \mid \eta)}{\partial \eta} d(\tilde{u}_i, \hat{\eta}_i).$$

This upper bound is continuous in π and achieves a maximum on the compact set $[\pi^\dagger, 1 - \pi^\dagger]$. The derivative is bounded above; a similar argument generates a lower bound. Writing B for the bound on the absolute value of the derivative, choose $\delta \leq \min \left\{ \pi^\dagger, \frac{\varepsilon}{B} \right\}$, and consider a pair π_L and π_H where $0 < \pi_H - \pi_L \leq \delta$. Now, if $\pi^\dagger \leq \pi_L < \pi_H \leq 1 - \pi^\dagger$

$F(p \mid \hat{\eta}) = \Pr[p(\pi) \leq p \mid \hat{\eta}]$. Next,

$$\lambda_n(\hat{\eta}) = \log \left[\frac{\int_0^1 p^{\bar{x}}(1-p)^{n-\bar{x}} dF(p \mid \hat{\eta})}{\int_0^1 (1-p)^{\bar{x}} p^{n-\bar{x}} dF(p \mid \hat{\eta})} \right],$$

where $F(p \mid \hat{\eta}) = \Pr[p(\pi) \leq p \mid \hat{\eta}]$. Given this, I form a best response by setting $\hat{v}(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} \geq 0]$ for $\tilde{u} \geq \bar{u}$, and for $|\tilde{u}_i| < \bar{u}$,

$$\hat{v}(\tilde{u}, \hat{\eta}) = \mathbb{I} \left[\log \left[\frac{U(\tilde{u})}{1-U(\tilde{u})} \right] + \lambda_n(\hat{\eta}) \geq 0 \right].$$

With this new strategy, I compute $\hat{p}(\pi) = E[\hat{v}(\tilde{u}, \hat{\eta}) \mid \pi]$. The final result is a continuous mapping from \bar{P} into P (and hence \bar{P}). A continuous mapping from a compact and convex subset of a Banach space has a fixed point, from Schauder's fixed-point theorem. Taking this fixed point $p^*(\pi)$, I can construct an associated Nash voting strategy in a finite electorate.

B.3. *The stability of a strategy-revision process*

Footnote 14 described a strategy-revision process, similar to the idea of fictitious play. Indexing a sequence of voting strategies by t , imagine that voters contemplate voting sincerely, so that $v_0(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} \geq 0]$. Anticipating sincere behaviour by others, a voter might then consider a naive (*i.e.* decision theoretic) strategy $v_1(\tilde{u}, \hat{\eta}) = \mathbb{I}[\tilde{u} + a_1 + b_1 \hat{\eta} \geq 0]$. Further anticipation leads to $v_2(\tilde{u}, \hat{\eta})$, and so on. Given the linearity of the voting strategies described in Lemma 3, this iterative best-response process is captured by $b_{t+1} = \hat{b}(b_t)$ and $a_{t+1} = \hat{a}(a_t, b_t)$.

Footnote 14 claimed that $b_t \rightarrow b^*$ as $t \rightarrow \infty$. Since $\hat{b}(b)$ is decreasing, and hence it is more convenient to consider the mapping $B(b) = \hat{b}^{(2)}(b) = \hat{b}(\hat{b}(b))$. Notice that $b^* = \hat{b}(b^*)$ is also a fixed point of B . Taking the derivative $B'(b) = \hat{b}'(\hat{b}(b))\hat{b}'(b)$, it follows that this is an increasing function, since $\hat{b}' \leq 0$. Consider a generic fixed point b , satisfying $B(b) = b$. Evaluate the derivative at this fixed point. This satisfies

$$B'(b) = b \left[\frac{\hat{b}\kappa^2 + \kappa^2}{\xi^2 + (\hat{b}^2 + 2\hat{b})\kappa^2} - \frac{1}{1+\hat{b}} \right] \times \hat{b} \left[\frac{b\kappa^2 + \kappa^2}{\xi^2 + (b^2 + 2b)\kappa^2} - \frac{1}{1+b} \right] \tag{39}$$

$$= \left[\frac{(\hat{b}^2 + \hat{b})\kappa^2}{\xi^2 + (\hat{b}^2 + 2\hat{b})\kappa^2} - \frac{\hat{b}}{1+\hat{b}} \right] \times \hat{b} \left[\frac{(b^2 + b)\kappa^2}{\xi^2 + (b^2 + 2b)\kappa^2} - \frac{b}{1+b} \right]. \tag{40}$$

Both terms are less than 1, and hence $B'(b) < 1$ at a fixed point. It follows that any fixed point must be a downcrossing. Further fixed points would require an upcrossing, and hence there is a unique fixed point b^* . From this it follows that $b_t \rightarrow b^*$. To see this, notice that $b_{t+2} = B(b_t)$. From the properties of B , there is the required convergence. It is now straightforward to observe that $a_t \rightarrow \infty$ so long as $\zeta > b^*$.

B.4. *The impact of strategic voting*

In Section 3.3 I “equipped” a voter with a signal realization $\hat{\eta}_i = \eta$ and then calculated $\hat{p} = \Phi((1+b)\Phi^{-1}(\pi))$. A slightly different expression arises when I condition on her actually receiving a signal realization $\hat{\eta}_i = \eta$. This is because her preferences form part of the signal, and hence, conditional on $\hat{\eta}_i = \eta$, it is no longer the case that $\tilde{u}_i \sim N(\eta, \xi^2)$. In fact, Bayesian updating following observation of $\hat{\eta}_i = \eta$ yields

$$\tilde{u}_i \mid \eta, \hat{\eta}_i = \eta \sim N \left(\eta, \frac{\xi^2(m-1)}{m} \right) \Rightarrow p = \Phi \left(\sqrt{\frac{m}{m-1}} \frac{(1+b)\eta}{\xi} \right),$$

then $|p(\pi_H) - p(\pi_L)| \leq \delta B \leq \varepsilon$. If $\pi_L < \pi^\dagger$ then $\pi_H \leq 2\pi^\dagger$, and $|p(\pi_H) - p(\pi_L)| \leq p(2\pi^\dagger) \leq \varepsilon$, by construction of π^\dagger and equation (38); a similar argument can be used for $\pi_H \geq 1 - \pi^\dagger$. This establishes equicontinuity. (Interestingly, to establish equicontinuity across the whole interval $[0, 1]$, the endpoints need to be considered separately, since the derivative of $p(\pi)$ can blow up for $\pi \in \{0, 1\}$. This separate consideration relies upon equation (38), which in turn stems from the assumption that extreme preference types have a weakly dominant strategy to vote sincerely.) Clearly, the set P is uniformly bounded. I now extend the set to its closure. Take a convergent sequence $\{p_t(\pi)\}$ from P . The limit of an equicontinuous sequence is continuous. Moreover, adding the limit to P does not destroy equicontinuity. To see this, for fixed $\varepsilon > 0$, choose δ so that any member of the set p changes by at most $\frac{\varepsilon}{2}$ over any interval of size smaller than δ . Going far enough down the sequence so that sequence members are closer than $\frac{\varepsilon}{4}$ to the limit everywhere (which is assured, since the sequence converges uniformly by the use of the sup norm) the limit can change by at most ε over an interval smaller than δ . Thus the closure \bar{P} is a closed, bounded, and equicontinuous set. Following the Ascoli–Arzelà theorem, a closed, bounded, and equicontinuous set of functions forms a compact subset of a Banach space. By inspection, the set P and its closure \bar{P} are convex. (For instance, for two different members of the set P and hence two different voting strategies, simply take a convex combination of the voting strategies and obtain a new member of P .)

which differs only slightly from equation (13). Turning to the political impact of strategic voting, equation (14) follows from the calculation

$$p = \Phi\left(\frac{(1+b)\eta}{\sqrt{\text{var}[b\hat{\eta}_i + \tilde{u}_i \mid \eta]}}\right) = \Phi\left((1+b)\Phi^{-1}(\pi)\sqrt{\frac{m}{(m-1)+(1+b)^2}}\right), \quad (41)$$

which can be inverted to obtain π in terms of p .

B.5. Specification of idiosyncrasy

One way to specify ξ^2 is to set two quantiles for voters with different preference intensities. For instance, a voter prefers candidate 1 when $u_1 \geq u_2$, and hence I may define $\pi(1) = \Pr[u_1 \geq u_2] = \Phi(\eta/\xi)$. A voter prefers candidate 1 k times as much as candidate 2 (relative to the status quo) when $u_1 \geq ku_2 \Leftrightarrow \tilde{u} \geq \log k$, yielding $\pi(k) = \Pr[u_1 \geq ku_2] = \Phi((\eta - \log k)/\xi)$. These two equations tie down η and ξ^2 :

$$\begin{aligned} \pi(k) &= \Phi((\eta - \log k)/\xi) \Rightarrow \xi \Phi^{-1}(\pi(k)) = \eta - \log k = \xi \Phi^{-1}(\pi(1)) - \log k \\ &\Rightarrow \xi = \frac{\log k}{\Phi^{-1}(\pi(1)) - \Phi^{-1}(\pi(k))}. \end{aligned}$$

To see this formula in action, consider an electorate where a fraction $\pi = \pi(1) = 0.6$ of the electorate rank candidate 1 highest, and half of these (or a fraction $\pi(2) = 0.3$ of the electorate) prefer candidate 1 twice as much as candidate 2. Then,

$$\xi = \frac{\log 2}{\Phi^{-1}(0.6) - \Phi^{-1}(0.3)} = 0.891 \quad \Rightarrow \quad \xi^2 = 0.794.$$

Using this formula, ξ^2 may also be specified by considering the median supporter of candidate 1. If the median supporter prefers candidate 1 k times as much as candidate 2, then (by definition of being the median in this group) $\pi(k) = \pi(1)/2$. Hence $\xi = \log k / [\Phi^{-1}(\pi) - \Phi^{-1}(\pi/2)]$. When $\pi = 1/2$ this formula reduces to $\xi = -\log k / \Phi^{-1}(1/4)$. Illustrating this last formulation, consider a balanced electorate in which $\pi = 1/2$. Suppose that the median supporter of candidate 1 prefers her favoured candidate twice as much as candidate 2. Then $\xi = 1.028$ and $\xi^2 = 1.056$. For $k = 1.5$ and $k = 2.5$ the outcomes are $\xi = 0.6$ and $\xi = 1.36$ respectively. Hence a range of $0.5 < \xi < 1.5$ might seem appropriate for the idiosyncrasy parameter.

B.6. Calibration of the U.K. General Election of 1997

Excluding the Speaker's seat and Tatton (where major parties were absent), for the other 527 English constituencies the leading candidates were the three major parties. Writing ψ_0 for the Conservative share among these three, I excluded $\psi_0 \geq 1/2$ (a certain Tory win) and $\psi_0 \leq 1/3$ (a certain Tory loss). For the remaining 270 constituencies I labelled the leading challenger as 1 and the trailing challenger as 2. This led to $\gamma = \psi_0/(1 - \psi_0)$ and $p = \psi_1/(\psi_1 + \psi_2)$. Using $m = 1/\omega$, I calculated the response parameter b . I then employed equation (14).

Acknowledgements. I thank friends, colleagues, seminar participants, the editor, and referees for their comments and assistance. Special thanks go to Steve Fisher, Justin Johnson, and Chris Wallace.

REFERENCES

- BAKER, L., KOESTNER, R., WORREN, N. M., LOSIER, G. F. and VALLERAND, R. J. (1995), "False Consensus Effects for the 1992 Canadian Referendum", *Canadian Journal of Behavioural Science*, **27** (2), 212–225.
- BROWN, C. E. (1982), "A False Consensus Bias in 1980 Presidential Preferences", *Journal of Social Psychology*, **118** (1), 137–138.
- CAIN, B. E. (1978), "Strategic Voting in Britain", *American Journal of Political Science*, **22** (3), 639–655.
- CARLSSON, H. and VAN DAMME, E. (1993), "Global Games and Equilibrium Selection", *Econometrica*, **61** (5), 989–1018.
- CHAMBERLAIN, G. and ROTHSCILD, M. (1981), "A Note on the Probability of Casting a Decisive Vote", *Journal of Economic Theory*, **25** (1), 152–162.
- COX, G. W. (1984), "Strategic Electoral Choice in Multi-Member Districts: Approval Voting in Practice?", *American Journal of Political Science*, **28** (4), 722–738.
- COX, G. W. (1994), "Strategic Voting Equilibria Under the Single Nontransferable Vote", *American Political Science Review*, **88** (3), 608–621.
- COX, G. W. (1997) *Making Votes Count* (Cambridge: Cambridge University Press).
- DEGROOT, M. H. (1970) *Optimal Statistical Decisions* (New York: McGraw-Hill).
- DROOP, H. R. (1871), "On the Political and Social Effects of Different Method of Electing Representatives", in *Papers, 1863–70*, ch. 3 (London: Juridical Society) 469–507.
- DUVERGER, M. (1954) *Political Parties: Their Organization and Activity in the Modern State* (New York: Wiley).

- EVANS, G., CURTICE, J. and NORRIS, P. (1998), "New Labour, New Tactical Voting? The Causes and Consequences of Tactical Voting in the 1997 General Election", in D. Denver, J. Fisher, P. Cowley and C. Pattie (eds.) *British Elections and Parties Review Volume 8: The 1997 General Election* (London and Portland, OR: Frank Cass) 65–79.
- EVANS, G. and HEATH, A. (1993), "A Tactical Error in the Analysis of Tactical Voting: A Response to Niemi, Whitten and Franklin", *British Journal of Political Science*, **23**, 131–137.
- FARQUHARSON, R. (1969) *Theory of Voting* (New Haven: Yale University Press).
- FEDDERSEN, T. J. and PESENDORFER, W. (1996), "The Swing Voter's Curse", *American Economic Review*, **86** (3), 408–424.
- FEDDERSEN, T. J. and PESENDORFER, W. (1997), "Voting Behavior and Information Aggregation in Elections with Private Information", *Econometrica*, **65** (5), 1029–1058.
- FEDDERSEN, T. J. and PESENDORFER, W. (1998), "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting", *American Political Science Review*, **92** (1), 23–35.
- FISHER, S. D. (2000), "Intuition versus Formal Theory: Tactical Voting in England 1987–1997" (Paper prepared for APSA Annual Meeting, Washington, DC).
- FISHER, S. D. (2001), "Extending the Rational Voter Theory of Tactical Voting" (Paper presented at the Mid-West Political Science Association Meeting, Chicago, IL).
- FISHER, S. D. (2005), "Definition and Measurement of Tactical Voting: The Role of Rational Choice", *British Journal of Political Science*, **34** (1), 152–166.
- FRANKLIN, M., NIEMI, R. G. and WHITTEN, G. (1994), "The Two Faces of Tactical Voting", *British Journal of Political Science*, **24**, 549–557.
- GOEREE, J. K. and GROSSER, J. (2003), "Costly Voting with Correlated Preferences" (Working Paper, University of Amsterdam).
- GOLDSTEIN, I. and PAUZNER, A. (2005), "Demand–Deposit Contracts and the Probability of Bank Runs", *Journal of Finance*, **60** (3), 1293.
- HEATH, A. and EVANS, G. (1994), "Tactical Voting: Concepts, Measurements and Findings", *British Journal of Political Science*, **24**, 557–561.
- HOFFMAN, D. T. (1982), "A Model for Strategic Voting", *SIAM Journal of Applied Mathematics*, **42**, 751–761.
- MCKELVEY, R. and ORDESHOOK, P. C. (1972), "A General Theory of the Calculus of Voting", in J. F. Herndon and J. L. Bernd (eds.) *Mathematical Applications in Political Science*, Vol. 6 (Charlottesville: University Press of Virginia) 32–78.
- MORRIS, S. and SHIN, H. S. (1998), "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks", *American Economic Review*, **88** (3), 587–597.
- MORRIS, S. and SHIN, H. S. (2003), "Global Games: Theory and Applications", in M. Dewatripont, L. Hansen and S. Turnovsky (eds.) *Advances in Economics and Econometrics, the Eighth World Congress* (Cambridge: Cambridge University Press) 56–114.
- MORRIS, S. and SHIN, H. S. (2004), "Coordination Risk and the Price of Debt", *European Economic Review*, **48** (1), 133–153.
- MORTON, R. B. (2006) *Analyzing Elections*, ch. 14 (New York: W. W. Norton).
- MYATT, D. P. and FISHER, S. D. (2002a), "Tactical Coordination in Plurality Electoral Systems", *Oxford Review of Economic Policy*, **18** (4), 504–522.
- MYATT, D. P. and FISHER, S. D. (2002b), "Everything is Uncertain and Uncertainty is Everything" (Department of Economics Discussion Paper Number 115, University of Oxford).
- MYERSON, R. B. and WEBER, R. J. (1993), "A Theory of Voting Equilibria", *American Political Science Review*, **87** (1), 102–114.
- NIEMI, R. (1984), "The Problem of Strategic Voting under Approval Voting", *American Political Science Review*, **78** (4), 952–958.
- NIEMI, R., WHITTEN, G. and FRANKLIN, M. (1992), "Constituency Characteristics, Individual Characteristics and Tactical Voting in the 1987 British General Election", *British Journal of Political Science*, **22**, 229–240.
- NIEMI, R. G., WHITTEN, G. and FRANKLIN, M. (1993), "People Who Live in Glass Houses: A Response to Evans and Heath's Critique of our Note on Tactical Voting", *British Journal of Political Science*, **23**, 549–563.
- PALFREY, T. R. (1989), "A Mathematical Proof of Duverger's Law", in P. C. Ordeshook (ed.) *Models of Strategic Choice in Politics* (Ann Arbor: University of Michigan Press) 69–92.
- PATTIE, C. and JOHNSTON, R. (1999), "Context, Conversation and Conviction: Social Networks and Voting at the 1992 British General Election", *Political Studies*, **47** (5), 877–889.
- RIKER, W. H. (1982a) *Liberalism Against Populism: A Confrontation Between the Theory of Democracy and the Theory of Social Choice* (Prospect Heights, IL: Waveland Press).
- RIKER, W. H. (1982b), "The Two-Party System and Duverger's Law: An Essay on the History of Political Science", *American Political Science Review*, **76** (4), 753–766.
- RIKER, W. H. and ORDESHOOK, P. C. (1968), "A Theory of the Calculus of Voting", *American Political Science Review*, **62** (1), 25–42.
- ROSS, L., GREENE, D. and HOUSE, P. (1977), "The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes", *Journal of Experimental Social Psychology*, **13** (3), 279–301.
- SIMON, H. A. (1954), "Bandwagon and Underdog Effects and the Possibility of Equilibrium Predictions", *Public Opinion Quarterly*, **18**, 245–253.