

SELF-ATTENUATING STRATEGIC VOTING

DAVID P. MYATT

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

September 1999

ABSTRACT. This is an analysis of strategic voting under qualified majority voting. Existing formal analyses of the plurality rule predict complete coordination of strategic voting: a strict interpretation of Duverger's Law. This conclusion is rejected. Unlike previous models, the popular support for each option is not commonly certain. Agents base their vote on both public and private signals of popular support. When private signals are the main source of information, the uniquely stable equilibrium entails only limited strategic voting and hence partial coordination. This is due to the surprising presence of negative feedback — strategic voting is a self-attenuating phenomenon. The theory leads to the conclusion that multi-candidate support in a plurality electoral system is perfectly consistent with rational voting behaviour.

Incomplete. See Nuffield College DP Series for Completed Version.

1. RETHINKING STRATEGIC VOTING

Duverger (1954) introduced his Law to political economy by noting that “the simple-majority single-ballot system favors the two-party system”. His aim was to evaluate the effect of voting systems on the structure and number of political parties. Duverger's writing envisaged an ongoing process involving both voters and political parties with bipartism as an eventual conclusion. More recent authors have offered a stricter version of Duverger's Law. The models of Cox (1994), Palfrey (1989) and Myerson and Weber (1993) predict strict bipartism as the outcome of any plurality rule election. Palfrey's (1989) “mathematical proof” claims that:

“... with instrumentally rational voters and fulfilled expectations, multicandidate contests under the plurality rule should result in only two candidates getting any votes.”

These authors consider a population of agents each casting a single vote, where the candidate with the largest number of votes wins. They claim that the uniquely stable equilibrium outcome

This paper is based on part III of Myatt (1999). Stephen D. Fisher inspired this work with his extensive empirical research on tactical voting, and with many hours of conversation on the topic. Grateful thanks are also due to David Firth, Geoff Evans, Iain McLean, Chris Wallace and seminar participants at Oxford, UCL, LSE and Edinburgh for helpful comments. Robin Mason is entirely responsible for any remaining errors, due to his nagging request to get the paper finished prior to a Southampton seminar. Cheers Robin.

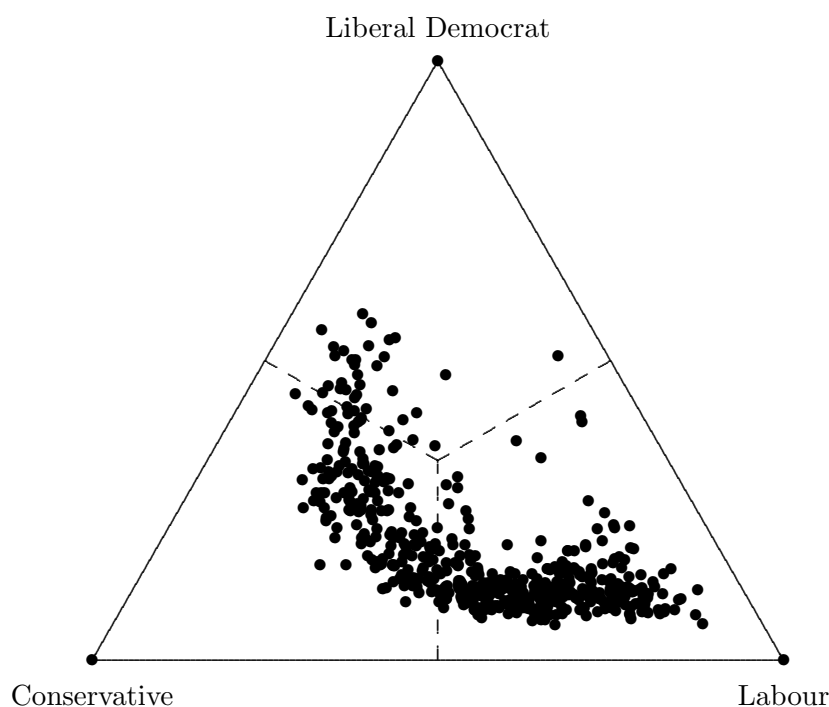


FIGURE 1. British General Election 1997. This simplex plot shows the relative vote shares for the three main British political parties — Conservative, Labour and Liberal Democrat. The plot restricts to English constituencies, hence avoiding the impact of the increased number of competing political parties in Wales, Scotland and Northern Ireland.

involves positive support for only two candidates. This outcome is the result of strategic voting — agents voting for other than their preferred candidates. Indeed, the prediction is that agents fully coordinate in their strategic behaviour.

Is the bipartite prediction borne out by the data? Unsurprisingly, it is not. The 1997 British General Election provides an illustration. Throughout the constituencies of England, three major parties compete. Figure 1 plots the relative vote shares for these three parties in 527 English parliamentary constituencies.¹ It appears that complete bipartite outcomes are absent. This appeal to the data would suggest a lack of rationality on voters — the degree of strategic voting is rather less than a rationality-based theory predicts. Or is it?

This new theory argues that partial coordination in a strategic voting setting is perfectly consistent with fully rational behaviour on the part of individual agents. The argument stems from the observation that the assumed independence of preferences drives existing models. In response, a new model is developed in which agents are uncertain of constituency-wide preferences. The public and private information sources upon which individuals condition their voting decisions are carefully specified. When private information dominates, the analysis shows that

¹Two English constituencies are omitted. West Bromwich West is the Speaker's seat, and is uncontested by major parties. The Labour and Liberal Democrat parties withdrew from Tatton in favour of an independent candidate. Thanks are due to Steve Fisher for providing the data.

strategic voting is self-attenuating rather than self-reinforcing. Negative feedback leads away from fully coordinated strategies, and towards multi-candidate support as a stable equilibrium.

The argument is built upon the analysis of a qualified majority voting game. This game is designed to highlight strategic vote switches between two candidates. Each member of a large electorate must cast a vote for one of two options. Both of these options are preferred by all agents to the status quo. Implementation of an option requires a qualified majority of the electorate to vote in favour: For $1 > \gamma > \frac{1}{2}$, a fraction γ of the electorate must vote in favour of an option in order to avoid the status quo. An immediate strategic incentive is present. An agent, preferring option 2, may instead decide to vote for option 1 in the belief that her vote is more likely to influence the election, and hence avoid the disliked default outcome.

What determines the voting decision of a rational agent in this scenario? A single agent can only influence the outcome when she has a *casting vote*. This *pivotal outcome* occurs when the vote total for one of the options is just equal to that required for a qualified majority. An extra vote will then implement that option rather than the status quo. The agent balances the relative probability of the two possible pivotal outcomes against her relative preference for the two options. It is the *pivotal likelihood ratio* that is the key determinant of an agent's voting decision. This key insight is clear from earlier decision-theoretic work by Hoffman (1982) and McKelvey and Ordeshook (1972), and is explored in a game-theoretic context by Palfrey (1989), Myerson and Weber (1993) and Cox (1994).² Indeed, this likelihood ratio provides the strategic incentive for a rational voter to abandon her preferred option.

Unfortunately, these earlier models all share a common feature. The preferences of individual agents are assumed to be drawn independently from a commonly known distribution. Why is this feature so critical to their results? Suppose that all remaining agents commit to voting straightforwardly — they vote for their preferred option. As the constituency size grows large, the *absolute* probability of a pivotal outcome falls to zero. More importantly, however, the *relative* probability — the pivotal likelihood ratio — diverges as the constituency grows large. This yields an unboundedly large strategic incentive for a rational voter. Such a voter will almost always choose to vote for the option with greater constituency-wide support. Of course, adopting a game-theoretic perspective, this effect is reinforced, and a fully-coordinated equilibrium in which all agents vote for a single option is realised.

The independence of preferences that is so key to these earlier models is an unattractive feature. Whereas the preferences of an individual are unknown to an observer, the *average* constituency wide preference is certain. As the constituency grows large, the idiosyncratic preferences of individual agents are averaged out — a consequence of the Law of Large Numbers. An observer can then give a precise prediction of the electoral outcome, even for truthful voting. The theory presented here abandons this feature. To achieve this, an individual agent's relative

²For an authoritative survey of this literature see Cox (1997).

payoff for the two options is decomposed into *common* and *idiosyncratic* effects. The common effect is shared by all agents, whereas the idiosyncratic component is distributed independently throughout the electorate. Importantly — and in contrast to the Cox-Palfrey framework — the common effect is unknown to members of the electorate. As the constituency grows large, the idiosyncratic effects are averaged out, but uncertainty over the common effect remains. It follows that the pivotal likelihood ratio — and hence the strategic incentive — remains finite. Importantly, this limiting pivotal likelihood ratio is driven *entirely* by uncertainty over the common effect.

With a finite pivotal likelihood ratio, strategic voting is incomplete and there is only partial coordination. Returning to a game theoretic perspective, however, the possibility of a fully coordinated equilibrium outcome remains. The standard logic of self-reinforcing strategic voting is as follows. The loss of support for the less favoured option from strategic switching enhances the strategic incentive to switch to the more favoured option. Strategic switching increases once more, yielding a further increase in the strategic incentive. This is a tale of positive feedback, yielding the “bandwagon effect” of Simon (1954).

This logic is flawed. In fact, strategic voting may exhibit *negative* feedback — a self-attenuating phenomenon. What argument supports this claim? First, note that the positive-feedback logic makes the implicit assumption that the most-favoured option is commonly known. If voting decisions are based on the privately observed information sources, then this assumption fails. It is clear that a voter’s information sources are important.

To investigate this issue, the information sources upon which votes base their decisions are modelled. Agents commonly observe a *public* signal of the common utility component — this formalises the idea of a publicly observed opinion poll. Each individual agent also observes a *private* signal. This feature reflects private interaction with other members of the constituency. Voting decisions are then contingent on the realisation of this private signal, as well as payoffs and the public signal. Perhaps surprisingly, if all other agents *increase* their response to their private signal, then the best response for a rational individual is to *reduce* her response in turn.

Why is this? Consider a constituency in which all agents vote straightforwardly. A relatively large constituency-wide lead for option 1 is required to achieve the qualified majority, and similarly for option 2. When computing the pivotal likelihood ratio, a rational agent compares two events that are relatively far apart. Switch now to a constituency in which voters respond strongly to their private signals. A much *smaller* lead for option 1 is sufficient to achieve the qualified majority. A small lead yields private signals in favour of option 1. These translate into strategic votes away from option 2, and hence to the required qualified majority. Similarly, only a small lead for option 2 is sufficient to do the same. These two pivotal events are now much closer, yielding a likelihood ratio that is closer to one. A rational vote faces a *lower* strategic incentive.

Of course, this argument relies on the use of private signals. If voting decisions are based on public signals, strategic voting continues to be self-reinforcing. Including both public and private signals, the analysis finds a unique partially coordinated equilibrium. This is uniquely stable when private signals are sufficiently precise relative to public signals. Hence, in constituencies where private information sources are likely to be more important, rational voting leads to a partially coordinated voting equilibrium.

Moving to a uniquely defined partially coordinated equilibrium allows a new range of comparative statics. Restricting to pure private information sources, the precision of private signals increases the incentive for tactical voting, although at a slower rate than in a decision-theoretic model. Tactical voting is also increasing in the severity of the qualified majority hurdle, as well as the asymmetry between the support of the two options.

The argument is formalised in the following sections. Section 2 describes the qualified majority voting game, preferences and signals. The importance of constituency uncertainty is highlighted in the analysis of Section 4. Sections 5 and 6 demonstrate the self-attenuating nature of strategic voting and investigate the existence and stability of equilibria. Section 8 concludes.

2. MODEL

The argument is formalised with reference to a simple qualified majority voting game.

2.1. Voting Rules. There are $n + 1$ agents, indexed by $i \in \{0, 1, \dots, n\}$. A collective decision is taken via qualified majority voting. There are three possible actions $j \in \{0, 1, 2\}$, where $j = 0$ represents the status quo. Each agent casts a single vote for either of the two options $j \in \{1, 2\}$.³ Denoting the vote totals for each of these options by x_1 and x_2 respectively, it follows that $x_1 + x_2 = n + 1$. Based on these votes, the action implemented is:

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

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

³The model extends easily to deal with the possibility of abstentions, corresponding to the option $j = 0$, since all agents strictly prefer both of the actions $j \in \{1, 2\}$ to the status quo.

2.2. Preferences. Payoffs are contingent only on the implemented action. The payoff u_{ij} is received by agent i when action j is implemented. All agents strictly prefer both actions $j \in \{1, 2\}$ to the status quo. This yields the payoff normalisation $u_{i0} = 0$ and hence $\min\{u_{i1}, u_{i2}\} > 0$. The relative preference for the two options varies throughout the population of agents. Section 3 demonstrates that the ratio $[u_{i1}/u_{i2}]$ is sufficient to determine an agent's decision.

Assumption 1. *The ratio $[u_{i1}/u_{i2}]$ is decomposed into two components as follows:*

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

where η is a common component to all agents. The idiosyncratic component ϵ_i is distributed independently across agents, with distribution $\epsilon_i \sim N(0, \xi^2)$.

An easy interpretation is that η represents population-wide factors affecting all agents. By contrast, ϵ_i represents the idiosyncratic preference of agent i . More generally, η is the expectation of $\log[u_{i1}/u_{i2}]$ conditional on all population information, generating the residual component ϵ_i . The parametric specification of ϵ_i is not critical to the argument. Imposing a normal distribution allows an easy microfoundation for the signal specifications described below. Notice that the variance term ξ^2 provides a measure of idiosyncrasy throughout the population.

An individual agent does not observe the decomposition of her preferences. In particular, the common utility component η is unknown. Beliefs about this component are generated following the receipt of informative signals, to which the model specification now turns.

2.3. Signals. Agents begin with a common and diffuse prior over η , with no knowledge of the common utility component. Information on η is then gleaned from two sources: *public* and *private* signals. The public signal models the publication of opinion polls and similar surveys. Commonly observed by all, it is equal to the true value of the common component, plus noise.

Assumption 2. *Agents commonly observe a public signal $\mu \sim N(\eta, \sigma^2)$.*

Following observation of this signal, and prior to the receipt of any private information, voters update to a common public posterior belief of $\eta \sim N(\mu, \sigma^2)$.⁴ Although not a formal feature of the model, a microfoundation underpins Assumption 2. Suppose that the preferences of m_μ randomly chosen individuals are publicly observed. Indexing by k :

$$(1) \quad \mu = \frac{1}{m_\mu} \sum_{k=1}^{m_\mu} \log \left[\frac{u_{k1}}{u_{k2}} \right] \sim N \left(\eta, \frac{\xi^2}{m_\mu} \right)$$

⁴More formally, endow all agents with a common prior of $\eta \sim N(\eta_0, \sigma_0^2)$, and Bayesian update following the observation of μ . Allowing $\sigma_0^2 \rightarrow \infty$ yields the public posterior $\eta \sim N(\mu, \sigma^2)$.

It is clear that a public signal with variance $\sigma^2 = \xi^2/m_\mu$ is equivalent to the public observation of m_μ individuals. Notice that the derivation of Equation (1) employs the distributional specification of the idiosyncratic component. For large m_μ , however, the Central Limit Theorem suggests the normal as an appropriate specification for the distribution of μ .

Viewed as an opinion poll, Assumption 2 provides a natural framework. In particular, the widespread publication of opinion polls is common during an election. For large m_μ , this leads to precise public information. In an election scenario, however, publicly observed polls tend to occur at the *national* level. Voting, however, will typically take place at a *regional* level. At the regional level, public opinion polls are rather less common.⁵ Any common regional component to preferences will remain uncertain. Agents do, however, have other sources of information available to them. In particular, a signal of constituency-wide candidate support may be obtained from the people with whom an individual interacts. The important characteristic of such information is that it leads to *private* signals.

Assumption 3. *Each agent i observes a private signal $\delta_i \sim N(\eta, \kappa^2)$. Conditional on η , this is independent across agents but may be correlated with the idiosyncratic component ϵ_i .*

Once again, a microfoundation is available. The signal δ_i corresponds to the *private* observation of m_δ randomly chosen individuals, with $\kappa^2 = \xi^2/m_\delta$. In particular, an agent's own payoffs are a signal of η . Hence, with $m_\delta = 1$, it follows that $\kappa^2 = \xi^2$. More generally, with this microfoundation, the variance of private signals is bounded above, with $\kappa^2 \leq \xi^2$. The inclusion of an agent's own preferences in the signal results in correlation between the signal and idiosyncratic utility component ϵ_i . For instance, in a sample of size $m_\delta > 1$:

$$\delta_i = \eta + \frac{1}{m_\delta} \left[\epsilon_i + \sum_{k \neq i} \epsilon_k \right] \quad \Rightarrow \quad \text{cov}(\delta_i, \epsilon_i) = \frac{\xi^2}{m_\delta} = \kappa^2$$

This feature is incorporated into the analysis, and extends easily to further correlation between the preferences of voter i and sampled individuals. Defining the correlation coefficient between δ_i and ϵ_i as ρ , the microfoundation presented here yields:

$$\rho \geq \frac{\kappa}{\xi} > 0$$

Following the observation of δ_i , a voter updates to obtain a private posterior belief.

Lemma 1. *The posterior belief of voter i over η satisfies:*

$$\eta \sim N \left(\frac{\kappa^2 \mu + \sigma^2 \delta_i}{\kappa^2 + \sigma^2}, \frac{\kappa^2 \sigma^2}{\kappa^2 + \sigma^2} \right)$$

Proof: Apply the standard Bayesian updating procedure — see DeGroot (1970). □

⁵Once again, the 1997 British General Election provides an example. Evans, Curtice and Norris (1998) note that 47 nationwide opinion polls were conducted during the election campaign. By contrast, only 29 polls were conducted in 26 different constituencies at a constituency level, out of a total of 659 constituencies.

The specification of private signals implicitly assumes that sampled individuals reveal their preferences truthfully. Furthermore, since the unconditional probability of influencing the election outcome is vanishingly small, it seems unlikely that individuals would find a costly information acquisition exercise to be worthwhile. The argument presented here accepts this latter critique. If a voter finds it too costly to conduct a private opinion poll, then the strategic manipulation of voting intentions by sampled individuals is no longer of relevance. The question of the private information source remains. It is envisioned that his information is accumulated over an extended period of time prior to an election, in the course of daily activity. It seems unlikely that a sampled individual would find response manipulation worthwhile over such a time frame.

How does the specification of this model differ from the Cox-Palfrey framework? Consider a symmetric strategy profile. With such a profile, the voting decision of an agent is contingent solely on the realised signals and payoffs. Conditional on η and μ , the private signals and payoffs are identically and independently distributed across voters. It follows that voting decisions inherit these properties. This indeed would yield the Cox-Palfrey model. Note, however, that η is unknown to any particular agent. From the agent's point of view, the voting decisions of the remaining electorate are *not* independent. This crucial difference is central to the argument.

3. VOTING BEHAVIOUR

Consider the behaviour of agent $i = 0$. She may only influence the outcome of the election if she is *pivotal*. A pivotal situation arises if, absent her vote, there is an exact tie. Among the remaining n agents $i \geq 1$, write x for the total number of votes cast for option 1. There are two possible pivotal scenarios. If $x = \gamma_n n$, then an additional vote will implement option 1 rather than the status quo. Similarly, if $n - x = \gamma_n n \Leftrightarrow x = (1 - \gamma_n)n$, then a single vote will tip the balance to option 2. Agent $i = 0$ has a *casting vote*. Conditioning on any information available to agent $i = 0$, consider the behaviour of the remaining agents. Write:

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

Hence q_1 and q_2 are the *pivotal probabilities* for options 1 and 2, in which one more vote is required to implement each of these options. Voting for option 1 will turn the status quo into the implementation of action 1 with probability q_1 , an expected payoff of $q_1 u_1$, relative to abstention. Similarly, A vote for option 2 has expected payoff $q_2 u_2$. Although the formal specification of the model rules out abstention, it is clear that some vote is optimal whenever $\min\{q_1, q_2\} > 0$. This argument leads to the following simple lemma.

Lemma 2. *For a voter with payoffs u_1 and u_2 , an optimal voting rule must satisfy:*

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

where q_j are the perceived pivotal probabilities.

Importantly, notice that when option 1 is strongly supported, so that $x > \gamma_n n$, an agent's vote has no effect. Similarly when option 2 is strongly supported. It follows that a rational voter is uninterested in such events, and concerned *only* with the probability of tied outcomes. This notion, familiar from earlier work by Hoffman (1982) and Myerson and Weber (1993) *inter alia*, will prove useful in developing intuition for the results that follow.

Before proceeding with the analysis, focus on the case where pivotal outcomes for both options are possible, so that $\min\{q_1, q_2\} > 0$. In this case, assume without loss of generality that an indifferent agent casts her vote for option 1.

Definition 1. Define the pivotal log likelihood ratio as $\lambda = \log[q_1/q_2]$.

Employing this definition, the optimal voting rule becomes:

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

It is clear that the key statistic of interest to a rational voter is λ , the pivotal log likelihood ratio. This is evaluated conditional on the appropriate strategy profile adopted by the remaining population. Indeed, a voting strategy for a rational agent may be conveniently characterised by the pivotal log likelihood ratio. This is formalised with the following definition.

Definition 2. A belief rule λ_i maps the signals of player i to the extended real line.

Using this definition, a rational agent supports candidate 1 whenever $\log[u_{i1}/u_{i2}] + \lambda_i \geq 0$. A belief rule of $\lambda_i \equiv +\infty$ corresponds to always voting for option 1, and symmetrically for $\lambda_i \equiv -\infty$. The analysis seeks symmetric Bayesian Nash equilibria of the qualified majority voting game, and hence the restriction will be to symmetric belief rules. Furthermore, rationality requires that an agent's beliefs (represented by λ_i) depend on her payoffs insofar as her payoffs yield relevant information. Given the microfoundation for private signals, all information provided by payoffs is reflected in δ_i . Indeed, μ and δ combine to yield a sufficient statistic for η . It follows that a restriction to beliefs rules that are contingent on δ_i and μ alone is without loss of generality.

Assumption 4. Restrict to symmetric belief rules, so that λ_i depends only on the signals, and not on the identity of player i . Dependence on μ will be suppressed, yielding $\lambda_i = \lambda(\delta_i)$.

Particular classes of belief rules will be of relevance. These are as follows.

Definition 3. A degenerate belief rule satisfies $\lambda \equiv \pm\infty$. A monotonic belief rule $\lambda(\delta_i)$ is strictly increasing, finite valued and differentiable in δ_i for all μ . An affine belief rule satisfies $\lambda(\delta_i) = a + b\delta_i$ for some a and b , where these parameters may depend on μ .

4. PIVOTAL PROPERTIES

4.1. **No Constituency Uncertainty.** Suppose that the common effect η is known. This is equivalent to the observation of a perfect public signal, where $\sigma^2 = 0$. Assumption 4 restricted to symmetric belief rules, and hence voting decisions are contingent solely on realised payoffs and signals. It follows that, conditional on η , these voting decisions are independent. Write p for the independent probability that agent i votes for option 1, so that:

$$p = \Pr \left[\log \left[\frac{u_{i1}}{u_{i2}} \right] + \lambda_i \geq 0 \mid \eta \right]$$

It follows that the vote total x for option 1 among the n agents $i \geq 1$ follows a binomial distribution with parameters p and n . The evaluation of the pivotal probabilities is straightforward:

$$q_1 = \Pr [x = \gamma_n n] = \binom{n}{\gamma_n n} p^{\gamma_n n} (1-p)^{(1-\gamma_n)n} \rightarrow 0 \quad \text{as } n \rightarrow \infty$$

Hence, for large constituencies, the *absolute* probability of a pivotal outcome falls to zero. But, as Section 3 argues, it is the *relative* likelihood of pivotal outcomes that drives the behaviour of a rational agent. Take the position of agent $i = 0$, and evaluate the pivotal log likelihood ratio:

$$\lambda_0 = \log \left[\frac{p^{\gamma_n} (1-p)^{1-\gamma_n}}{p^{1-\gamma_n} (1-p)^{\gamma_n}} \right]^n \rightarrow \begin{cases} +\infty & p > 1/2 \\ -\infty & p < 1/2 \end{cases}$$

It follows that for an unboundedly large electorate n , the tactical incentive λ_0 for agent $i = 0$ grows without bound. For $p > 1/2$, this agent will almost always choose to vote for option 1. Extending this response to the whole population, there is complete coordination, with all agents strategically abandoning option 2. Notice that the pivotal log likelihood ratio is entirely driven by idiosyncratic uncertainty. With independent preferences and signals, the consequent unbounded likelihood ratio drives strategic voting.

4.2. **Uncertain Common Effect.** But what if the common effect η is uncertain? Conditional on η , voting decisions continue to be drawn from a binomial distribution. But since p depends on η , and η is unknown, it follows that p is unknown. Consider, from the perspective of the focal agent, uncertainty over p and represent this by the density $f(p)$. Impose the following assumption on this density.

Assumption 5. *The density $f(p)$ is continuous and strictly positive on $(0, 1)$.*

Assumption 5 will be satisfied by the monotonic belief rules that are the focus of Section 5. The pivotal likelihood ratio becomes:

$$(2) \quad \frac{q_1}{q_2} = \frac{\int_0^1 [p^{\gamma_n} (1-p)^{1-\gamma_n}]^n f(p) dp}{\int_0^1 [p^{1-\gamma_n} (1-p)^{\gamma_n}]^n f(p) dp}$$

Now, as $n \rightarrow \infty$, both the numerator and denominator of this expression vanish. The ratio, however, does not diverge and indeed converges to an attractive expression. This is revealed in the following key proposition.

Proposition 1. *Allowing the electorate size to grow large:*

$$\lim_{n \rightarrow \infty} \frac{q_1}{q_2} = \frac{f(\gamma)}{f(1 - \gamma)}$$

Sketch Proof: Examining Equation (2), notice that the integrand in the numerator vanishes to zero due to the leading term $[p^\gamma(1 - p)^{1-\gamma}]^n$. As n grows, this expression becomes peaked around its maximum, which occurs at $p = \gamma$. Hence only density at this point contributes any weight. The denominator peaks at $p = 1 - \gamma$. At these points the leading terms of the integrands in the numerator and denominator coincide (after a simple change of variable) and hence the limiting likelihood ratio is the desired expression. See Appendix for a full proof. \square

What is the interpretation of this proposition? For large constituencies, the relative likelihood of ties involving options 1 and 2 is the relative likelihood that their respective constituency-wide support levels (represented by p) coincide with the critical value γ . Importantly, then, it is only uncertainty over p (generated by uncertainty over the common effect η) that matters. Why is this? As the constituency grows large, the individual idiosyncratic effects ϵ_i are averaged out. The common effect, however, cannot be averaged out and hence becomes the key determinant. Notice that in earlier models only idiosyncratic uncertainty was present. But in the presence of constituency uncertainty, its effect disappears. This suggests that the results from the Cox-Palfrey may be somewhat misleading in guiding our analysis of voting problems.

The second feature of the proposition is this: the limiting pivotal likelihood ratio is *finite*. This yields a finite strategic incentive. It follows that, if all other voters act straightforwardly, the tactical voting of the focal agent will not be complete. This leaves upon the possibility that tactical voting may be self-reinforcing, leading to an equilibrium outcome that involves full coordination. Addressing this issue requires a game-theoretic perspective, to which the analysis now turns.

5. BEST RESPONSE AND NEGATIVE FEEDBACK

5.1. Affine Belief Rules under Best Response. How does a rational voter incorporate her private signal? The affine belief rules of Definition 3 are attractive. In particular, notice that straightforward (truthful) voting corresponds to an affine belief rule with $a = b = 0$. Furthermore, with an extension to the real line, the degenerate belief rules (and hence fully coordinated strategy profiles) of $\lambda \equiv \pm\infty$ are obtained by setting $a = \pm\infty$ and $b = 0$. All well and good, but why should a rational agent use them? Attention is justified by the following.

Lemma 3. *Suppose that in an unboundedly large electorate, all agents $i \geq 1$ employ a monotonic belief rule $\lambda(\delta_i)$. A unique best response for agent $i = 0$ is to use an affine belief rule.*

Proof. When belief rules are monotonic, there are two unique values of η that result in a pivotal outcome. Posterior private beliefs over η are normal, and the log likelihood ratio of the normal is affine in its mean. For a formal proof, see Appendix. \square

A simple corollary is that the class of affine belief rules is closed under best response. Furthermore, any search for partially coordinated equilibria using monotonic belief rules may restrict to the affine class without loss of generality. The specific nature of best responses within the affine class is established in the next lemma.

Lemma 4. *The class of affine belief rules is closed under best response. If all adopt a belief rule $\lambda(\delta_i) = a + b\delta_i$, then a best response is to adopt an affine belief rule $\hat{\lambda}(\delta_i) = \hat{a} + \hat{b}\delta_i$ where:*

$$\begin{aligned}\hat{a} = \hat{a}(a, b) &= \frac{\hat{b}[a(\kappa^2 + \sigma^2) + (1 + b)\kappa^2\mu]}{\sigma^2(1 + b)} \\ \hat{b} = \hat{b}(b) &= \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)}{\kappa^2(1 + b)} \quad \text{where} \quad \tilde{\kappa}^2 = \text{var}[\epsilon_i + b(\delta_i - \eta) \mid \eta]\end{aligned}$$

Proof. See Appendix. \square

5.2. Self-Attenuating Strategic Voting. Section 1 argued that, with private signals as the dominant information source, strategic voting exhibits negative feedback. The best response function $\hat{b}(b)$ characterises an agent's response to her private signal. It's behaviour is described in the following lemma.

Lemma 5. *The mapping $\hat{b}(b)$ is strictly decreasing in b , with:*

$$\hat{b}(0) = \frac{2\xi\Phi^{-1}(\gamma)}{\kappa^2} \quad \text{and} \quad \lim_{b \rightarrow \infty} \hat{b}(b) = \frac{2\Phi^{-1}(\gamma)}{\kappa}$$

Proof. See Appendix. \square

Lemma 4 states that $\hat{b}(b)$ is *decreasing* in b : An increase in the tendency by others to vote strategically ($b \uparrow$) reduces the tendency for the rational agent $i = 0$ to vote strategically ($\hat{b} \downarrow$). Of course, the coefficient b measures only the response of the strategic incentive to the private signal. Turning to the intercept a , notice that:

$$\hat{a} = \frac{\hat{b}a(\kappa^2 + \sigma^2)}{\sigma^2(1 + b)} + \frac{\hat{b}\kappa^2\mu}{\sigma^2}$$

This consists of two terms. The first is a general strategic incentive term, and the second is the response to the public signal. Notice that *both* of these terms are decreasing in b . Fixing b , however, the intercept term is increasing in a . It follows that any common strategic incentive is self-reinforcing in the standard way.

What explains the presence of negative feedback? Intuition is aided by Figures 2 and 3. Begin by setting $a = b = 0$ for all agents $i \geq 1$, corresponding to straightforward voting. Conditional on the common component, a randomly selected agent supports option 1 with probability

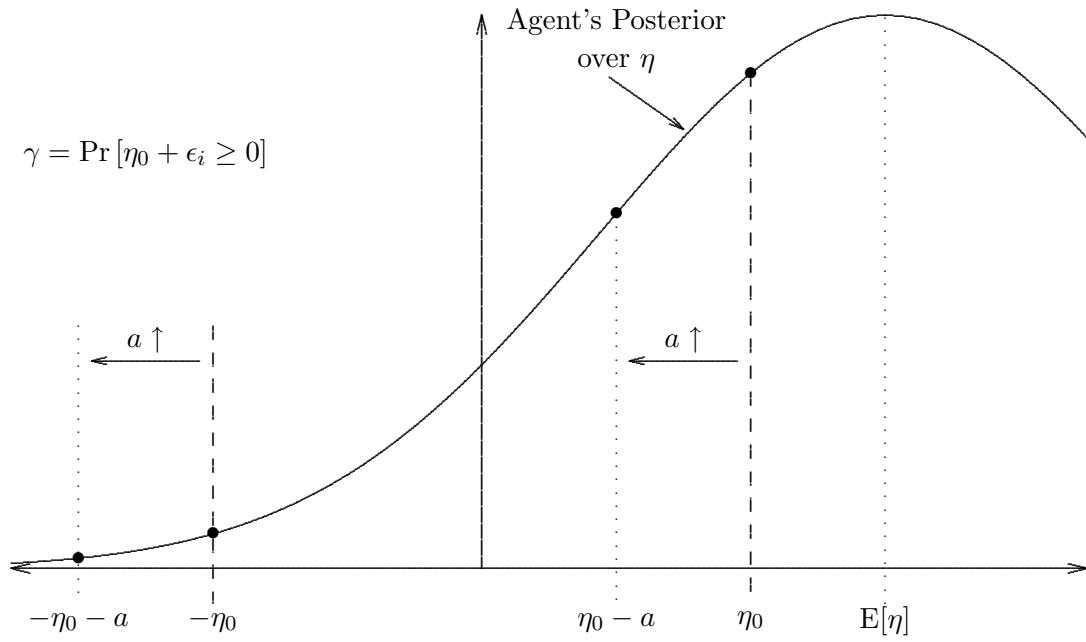


FIGURE 2. Self-Reinforcing Strategic Voting with Public Signals

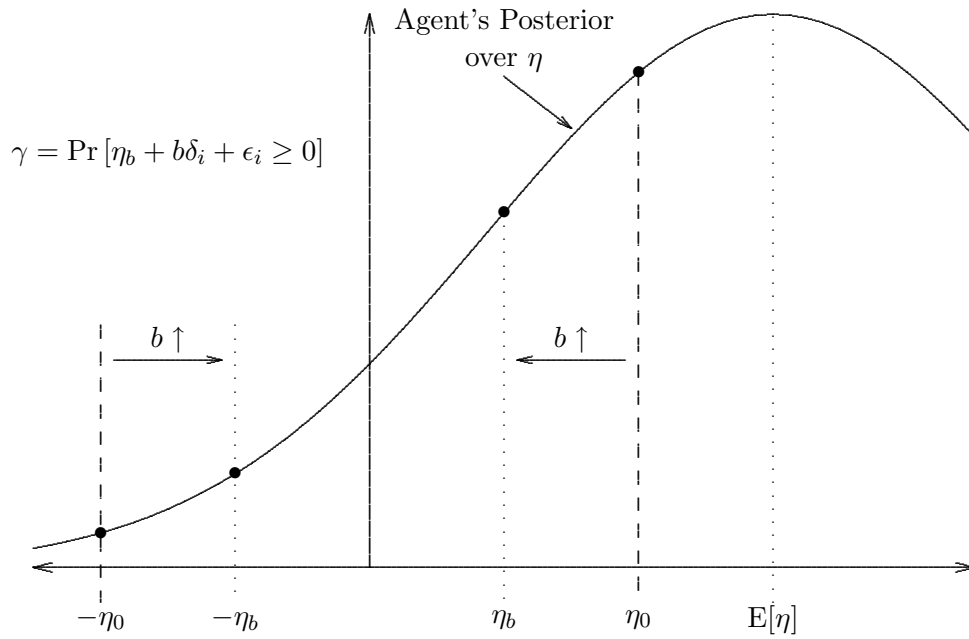


FIGURE 3. Self-Attenuating Strategic Voting with Private Signals

$p = \Pr[\eta + \epsilon_i \geq 0]$. The rational agent $i = 0$ computes the likelihood ratio of $p = \gamma$ versus $p = 1 - \gamma$. The critical values of the common component are η_0 and $-\eta_0$ respectively, where $\gamma = \Pr[\eta_0 + \epsilon_i \geq 0]$. These critical values are illustrated in Figures 2 and 3.

Now, increase a , so that $a > b = 0$. The critical values of the common component are now $\eta_0 - a$ and $-\eta_0 - a$. Comparing these two yields a larger likelihood ratio, and hence an increased strategic incentive. This is the standard logic of positive feedback — see Figure 2.

Resetting $a = 0$, now increase b , so that $b > a = 0$. The critical value of the common component for $p = \gamma$ is now η_b , where $\gamma = \Pr[\eta + b\delta_i + \epsilon_i \geq 0]$. Recalling that δ_i is symmetrically distributed around η , it follows that the critical value for $p = 1 - \gamma$ is $-\eta_b$. Notice that:

$$-\eta_0 < -\eta_b < 0 < \eta_b < \eta_0$$

Agent $i = 0$ computes the likelihood ratio for two common component values that are *closer* together. This yields a likelihood ratio that is closer to one, and hence a *reduced* strategic incentive. This effect is illustrated in Figure 3 — there is negative feedback.

This may at first seem counter-intuitive. When b is high, agents respond strongly to their signals. In particular, this increases the likelihood of a tactical vote. Importantly, it increases the probability of a strategic vote in *both* directions. Agent $i = 0$ with signal $\delta_0 > 0$ is concerned that other agents may observe signals $\delta_i < 0$, yielding a pivotal outcome involving option 2. For high δ_0 , this event seems most unlikely — surely option 1 will almost certainly win? But if option 1 will almost certainly win, then the vote of agent $i = 0$ has no effect. She can *only* influence the outcome when there is a tie. But if there is a tie, then her strong signal must have overstated the true constituency-wide support of option 1. She must therefore envisage a much lower true value for η . It is then reasonable for the vote to consider true values the common component satisfying $\eta < 0$.

6. EQUILIBRIUM

6.1. Full Coordination. Fully coordinated equilibria are always present in this model. To see this, suppose that all agents cast their vote for for option 1, irrespective of the public signal, their payoffs or the realisation of their private signals. It follows that $x = n$, and thus $q_2 = 0$. It is (weakly) optimal for agent $i = 0$ to retain the posited strategy profile. A symmetric argument establishes an equilibrium in which all agents vote for option 2.

Proposition 2. *There are two fully coordinated equilibria, where all vote for one option.*

Such equilibria correspond to pivotal log likelihood ratios of $\lambda = \pm\infty$.

6.2. Partial Coordination. A partially coordinated equilibrium is also available. As argued in Section 5, a restriction to monotonic belief rules yields a best response in the class of affine belief rules. An equilibrium in this class then corresponds to a finite pair $\{a^*, b^*\}$ such that $b^* = \hat{b}(b^*)$ and $a^* = \hat{a}(a^*, b^*)$. The properties of the mapping $\hat{b}(b)$ immediately yield a unique fixed point.

Lemma 6. *The mapping $\hat{b}(b)$ has a unique fixed point $b^* > 0$. For $\rho \geq \kappa/\xi$, this satisfies:*

$$(3) \quad \frac{2\Phi^{-1}(\gamma)}{\kappa} \leq b^* \leq \frac{\Phi^{-1}(\gamma)}{\kappa} \left\{ 1 + \sqrt{\frac{\Phi^{-1}(\gamma) + 2\xi}{\Phi^{-1}(\gamma)}} \right\}$$

For the microfoundation case, with $\rho = \kappa/\xi$, the bound may be refined to:

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

Proof. Uniqueness follows from Lemma 5. See Appendix for derivation of the inequalities. \square

Any partially coordinated equilibrium must entail $b = b^*$. It remains to consider fixed points of \hat{a} . Notice that \hat{a} is affine in a , yielding a unique partially coordinated equilibrium.

Proposition 3. *For $\sigma^2 \neq b^*\kappa^2$ there is a unique affine equilibrium satisfying:*

$$a^* = \frac{b^*(1 + b^*)\kappa^2\mu}{\sigma^2 - b^*\kappa^2}$$

Proof. Straightforward solution to $a^* = \hat{a}(a^*, b^*)$. \square

Corollary 1. *There is a unique equilibrium monotonic belief rule.*

An equilibrium selection problem is posed. This is resolved in the next section.

6.3. Equilibrium Selection. Begin with the initial hypothesis that all agents vote straightforwardly for their preferred option. This is equivalent to employing an affine belief rule with parameters $a_0 = b_0 = 0$. Agent $i = 0$, acting optimally in response to this strategy profile will employ an affine belief rule with parameters $a_1 = \hat{a}(0, 0)$ and $b_1 = \hat{b}(0)$. Of course, this agent will anticipate a similar response by the population at large, and hence update once more to obtain a belief rule of $a_2 = \hat{a}(a_1, b_1)$ and $b_2 = \hat{b}(b_1)$. This thought experiment describes an iterative best response process within the class of affine belief rules. Of course, a starting point within the monotonic class will enter and remain affine within one step. Formally:

Definition 4. *Define the iterative best response process by $b_t = \hat{b}(b_{t-1})$, $a_t = \hat{a}(a_{t-1}, b_{t-1})$.*

Having defined this process, global stability may be used as an appropriate equilibrium selection criterion. Begin with the mapping $\hat{b}(b)$. This mapping and the associated process $\{b_t\}$ are not contingent on a_t , and hence may be considered in isolation.

Lemma 7. b^* is globally stable in the iterative best response dynamic: $b_t \rightarrow b^*$ as $t \rightarrow \infty$.

Proof. See Section Appendix. □

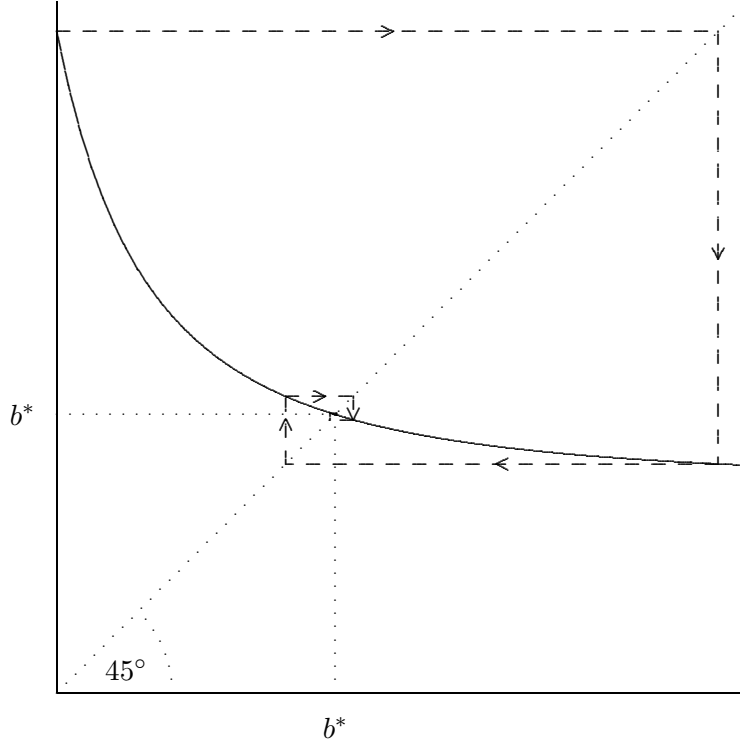


FIGURE 4. $\hat{b}(b)$ and Convergence to Equilibrium

Whereas the formal proof of Lemma 7 is algebraically tedious, a diagrammatic illustration proves useful. Figure 4 plots the best response function $\hat{b}(b)$, illustrating the convergence to the fixed point. Notice the cyclic behaviour — this is a consequence of the negative feedback inherent in strategic voting with private signals. Begin with $b_0 = 0$. Taking the next step, the rational agent recognises the strategic behaviour of others. This attenuates the response to the private signal, with a consequent reduction in b . Of course, this behaviour leaves open the possibility of a limit cycle in the iterative best response process. Lemma 7 ensures that the cycle dampens down, eventually converging to the unique fixed point b^* .

To select an equilibrium, turn to the mapping $\hat{a}(a, b^*)$.

Proposition 4. *If $\sigma^2 > b^* \kappa^2$, then a^* is uniquely stable. A sufficient condition is for:*

$$\sigma^2 > \kappa \Phi^{-1}(\gamma) \left\{ 1 + \sqrt{\frac{\Phi^{-1}(\gamma) + 2\xi}{\Phi^{-1}(\gamma)}} \right\}$$

If this holds, then the partially coordinated equilibrium is uniquely stable, and attained as the limit of the iterative best response process from any finite starting point.

Proof. This follows from the affine nature of the mapping \hat{a} — see Lemma 4. The sufficient condition is obtained by employing the upper bound on b^* from Lemma 6. \square

Corollary 2. *As $\sigma^2 \uparrow \infty$, the partially coordinated equilibrium is uniquely stable, with $a^* \downarrow 0$.*

It follows that the partially coordinated equilibrium is selected whenever the public information source is sufficiently imprecise. Notice that the precision is judged *relative* to the precision of the private signal.⁶ Hence if private signals are relatively more important than public signals, a partially coordinated outcome emerges. As Corollary 2 confirms, with only private information, the equilibrium is *never* fully coordinated. Which situation is likely to obtain? In a national referendum or similarly nationally conducted election, there are typically many public information sources. Moreover, region-specific effects are unimportant. It follows that public sources are likely to be more important than private. At a district level, however, commonly observed public signals are likely to be fewer. It follows that private signals are the primary source of information. A partially coordinated conclusion emerges, with multi-candidate support.

6.4. Comparative Statics. Comparative statics are absent in the Cox-Palfrey model of plurality voting. The prediction of complete strategic voting yields a prediction that is unresponsive to parameter changes. This is not the case here. This section explores comparative statics for the pure private information case, where $\sigma^2 \uparrow \infty$ and $a^* = 0$. Part II of Myatt (1999) provides an analysis of the pure public information case.

Recall from Section 3 that λ_i provides an appropriate measure of the strategic incentive for a voter. Without loss of generality, set $\eta > 0$ so that option 1 is preferred by the majority of the electorate. The expected strategic incentive is then $E[\lambda_i] = E[b^*\delta_i] = b^*\eta$. Employing Lemma 6, obtain:

$$\frac{2\Phi^{-1}(\gamma)\eta}{\kappa} \leq E[\lambda_i] \leq \frac{2\Phi^{-1}(\gamma)\eta}{\kappa} \sqrt{2 + 2\sqrt{\frac{(\Phi^{-1}(\gamma))^2 + \xi^2}{(\Phi^{-1}(\gamma))^2}}}$$

where the microfoundation case of $\rho = \kappa/\xi$ has been imposed. It is immediate that this incentive is increasing in qualified majority required (γ), as well as the precision of the private signal ($1/\kappa^2$). Care is required in varying η and ξ^2 , since this alters the true support for option 2. Denote the (expected) proportion of the population in favour of option 1 as p . Recall that $p = \Phi(\eta/\xi)$. It follows that $\eta = \xi\Phi^{-1}(p)$. Next, recall from the private signal microfoundation that $\kappa^2 = \xi^2/m_\delta$, where m_δ is the size an agent's private sample. The bounds become:

$$2\sqrt{m_\delta}\Phi^{-1}(\gamma)\Phi^{-1}(p) \leq E[\lambda_i] \leq 2\sqrt{m_\delta}\Phi^{-1}(\gamma)\Phi^{-1}(p) \sqrt{2 + 2\sqrt{\frac{(\Phi^{-1}(\gamma))^2 + \xi^2}{(\Phi^{-1}(\gamma))^2}}}$$

⁶In fact, the condition presented here is of the same form as that in Morris and Shin (1999). They consider the coordination problem of bank runs, and find that the ratio of the variance of a public signal and standard deviation of a private signal determines the uniqueness of equilibrium.

Clearly, the strategic incentive is increasing in p . Hence an increase in the (expected) asymmetry between the two options increases the incentive for supporters of the weaker option to shift their vote. Furthermore, the idiosyncrasy of preferences (ξ) increases the strategic incentive.

It is worthwhile considering the decision-theoretic benchmark — the behaviour of a rational agent in an electorate of straightforwardly-voting agents. To do this, evaluate the tactical incentive at the parameter $b_1 = \hat{b}(0)$. From the definition of \hat{b} from Lemma 4:

$$b_1 = \hat{b}(0) = \frac{2\xi\Phi^{-1}(\gamma)}{\kappa^2}$$

Evaluation the tactical incentive:

$$E[\lambda_i] = b_1\eta = 2m_\delta\Phi^{-1}(\gamma)\Phi^{-1}(p)$$

Notice the crucial role played by the precision of private information. In the decision-theoretic case the tactical incentive increases with the precision m_δ . By contrast, in the game-theoretic case, it increases with the *square root* of the precision — a much slower rate. Importantly, this reflects the self-attenuating nature of strategic voting. In an informal way, adding sophistication to agent's behaviour — moving from a decision-theoretic to a game-theoretic model — *reduces* the incentive for, and hence amount of, strategic voting.

The comparative static analysis presented above considers the strategic *incentive* for voters. Of course, the *probability* of observing a strategic vote is further affected by the idiosyncrasy of individual agents. As idiosyncrasy increases ($\xi^2 \uparrow$) there are a greater number of extreme agents, who require a larger incentive to switch their vote. This suggests that an increase in idiosyncrasy will reduce the probability of a strategic vote. Note, however, that the strategic incentive is increasing in ξ . The two effects must be considered together.

To investigate this issue, adopt the microfoundation example for the private signal, so that δ_i is based on a private sample of m_δ agents, including the agent herself. Furthermore, suppose that in the case of the $m_\delta - 1$ other agents, the signal is accurate so that $\sum_{k \neq i} \epsilon_k = 0$ — this helps simplify the analysis. It follows that:

$$\eta + b\delta_i + \epsilon_i = (1 + b)\eta + \left[1 + \frac{b}{m_\delta}\right] \epsilon_i$$

The agent votes for option 1 whenever:

$$(1 + b)\eta + \frac{b + m_\delta}{m_\delta} \epsilon_i \geq 0 \quad \Leftrightarrow \quad \frac{\epsilon_i}{\xi} \geq -\frac{(1 + b)m_\delta\eta}{(b + m_\delta)\xi}$$

By contrast, the agent actually prefers option 1 whenever $\epsilon_i/\xi \geq -\eta/\xi$. It follows that such an agent votes strategically whenever:

$$-\frac{\eta}{\xi} > \frac{\epsilon_i}{\xi} \geq -\frac{(1 + b)m_\delta\eta}{(b + m_\delta)\xi}$$

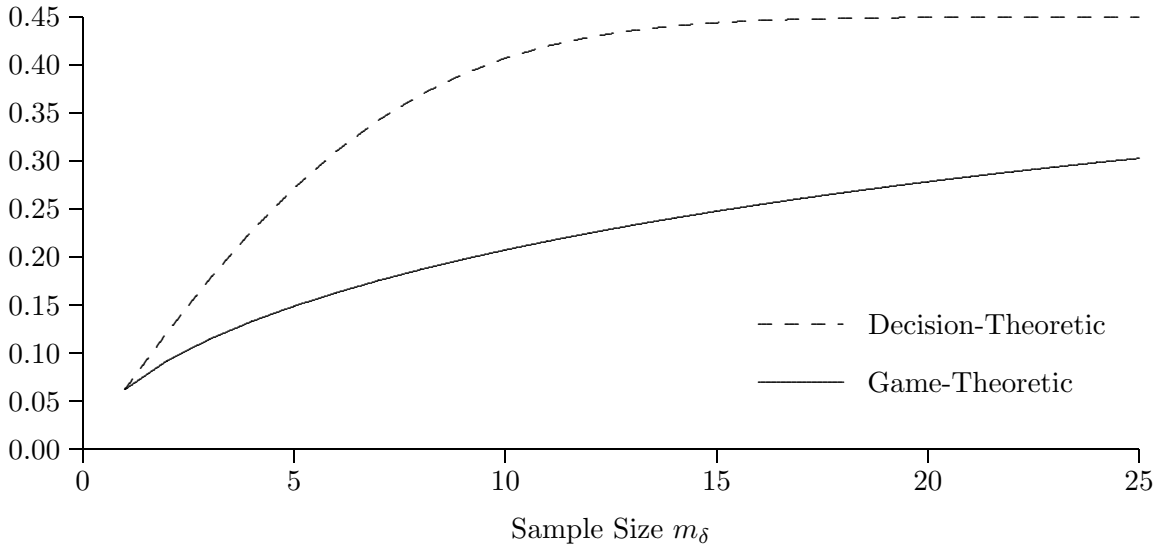


FIGURE 5. Strategic Voting Probabilities

Recalling that the true support for option $p = \Phi(\eta/\xi)$, the probability of a strategic vote is:

$$\Phi\left(\frac{(1+b)m_\delta}{b+m_\delta}\Phi^{-1}(p)\right) - p < \Phi((1+b)\Phi^{-1}(p)) - p$$

Comparative statics on the probability of a strategic vote are thus obtained via an analysis of b . This response b is actually decreasing in ξ , and hence an increase in the idiosyncrasy of agents' preferences will reduce the probability of a strategic vote.

7. ILLUSTRATION

A brief illustration of the results is helpful. Consider the following specification:

$$\begin{aligned}\gamma &= 0.6 \\ p &= 0.55 \\ \xi &= 0.4\end{aligned}$$

The parameter ξ is chosen so that approximately 2.5% of the population prefer option 2 twice as much as option 1. Using these parameters, Figure 5 plots the probability of observing a strategic vote against the precision of the information source. Both the game-theoretic case ($b = b^*$) and the decision-theoretic case ($b = \hat{b}(0)$) are displayed.

8. CONCLUSION

This new theory of strategic voting observes that the Cox (1994) and Palfrey (1989) models of strategic voting are driven by the assumption that voter preferences are drawn independently from a commonly-known distribution. This leads to the divergence of the pivotal log-likelihood ratios that are the critical determinants of optimal voting behaviour. Introducing uncertain

common effects to voter preferences results in uncertain constituency-wide candidate support, and overturns this result — pivotal log-likelihood ratios remain finite. Moreover, it is *only* uncertainty over common effects that matters. This suggests that the Cox-Palfrey model is perhaps driven by the wrong factors.

The introduction of uncertain common effects, and the modelling of voter information sources leads to new insights. The key rôle of information is now clear — a fully coordinated outcome requires both precise public information and the absence of precise private information. When regional effects are important, and information is likely to be privately observed, the model predicts some, but not complete, strategic voting.

Earlier models lack comparative static predictions. This is a necessary consequence of the strictly coordinated prediction. Here, the comparative statics are clear. Importantly, in a “close” election (corresponding to low γ and p), there is *less* strategic voting. Perhaps most importantly, strategic voting is self-attenuating, and increasing the sophistication of voters reduces its effect.

Of course, the formal model here is one of qualified majority voting. Myatt (1999) employs a variant of this model to address directly the issues arising in a plurality rule election. Weaknesses in that model (as here) remain. Indeed, uncertainty of the qualified majority (γ here) and multi-directional strategic voting require further analysis.

All of these issues are the subject of ongoing research. In addition, empirical testing of the model is already in progress, and experimental work is planned. Hopefully a greater understanding of strategic behaviour may lead to a better understanding of electoral systems.

APPENDIX A. OMITTED PROOFS

A.1. Pivotal Properties. This section provides omitted proofs from Section 4.

Proof of Proposition 1: Introduce the parameter $\tilde{\gamma}$ where $\frac{1}{2} < \tilde{\gamma} < 1$, and define $r(\tilde{\gamma})$ as follows:

$$r(\tilde{\gamma}) = \frac{\int_0^1 [p^{\tilde{\gamma}}(1-p)^{1-\tilde{\gamma}}]^n f(p) dp}{\int_0^1 [(1-p)^{\tilde{\gamma}} p^{1-\tilde{\gamma}}]^n f(p) dp} = \frac{\int_0^1 [p^{\tilde{\gamma}}(1-p)^{1-\tilde{\gamma}}]^n f(p) dp}{\int_0^1 [p^{\tilde{\gamma}}(1-p)^{1-\tilde{\gamma}}]^n f(1-p) dp}$$

where the second equality follows from a simple change of variables in the denominator. Recalling the definition of the pivotal probabilities q_1 and q_2 , it follows that $q_1/q_2 = r(\tilde{\gamma})$. Now introduce the notation $G(p)$:

$$(4) \quad G(p) \equiv \frac{p^{\tilde{\gamma}}(1-p)^{1-\tilde{\gamma}}}{\tilde{\gamma}^{\tilde{\gamma}}(1-\tilde{\gamma})^{1-\tilde{\gamma}}} \quad \Rightarrow \quad r(\tilde{\gamma}) = \frac{\int_0^1 G(p)^n f(p) dp}{\int_0^1 G(p)^n f(1-p) dp}$$

Notice that $G(p)$ is increasing from $G(0) = 0$, attaining a maximum of $G(\tilde{\gamma}) = 1$ at $p = \tilde{\gamma}$, and then declining back to $G(1) = 0$. Next fix $\min\{1/4, 1 - \tilde{\gamma}\} > \epsilon > 0$. For notational convenience,

define:

$$\begin{aligned} f_{L,\epsilon}(x) &= \inf_{x-\epsilon \leq p \leq x+\epsilon} f(p) = \min_{x-\epsilon \leq p \leq x+\epsilon} f(p) > 0 \\ f_{H,\epsilon}(x) &= \sup_{x-\epsilon \leq p \leq x+\epsilon} f(p) = \max_{x-\epsilon \leq p \leq x+\epsilon} f(p) < \infty \end{aligned}$$

where the extrema are well defined since $[x - \epsilon, x + \epsilon]$ is a compact set and $f(p)$ is continuous from Assumption 5. Employing this notation, formulate an upper bound for the ratio $r(\tilde{\gamma})$ in Equation (4):

$$(5) \quad r(\tilde{\gamma}) \leq \frac{f_{H,\epsilon}(\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp + f_{H,2\epsilon}(\tilde{\gamma}) \left[\int_{\tilde{\gamma}-2\epsilon}^{\tilde{\gamma}-\epsilon} G(p)^n dp + \int_{\tilde{\gamma}+\epsilon}^{\tilde{\gamma}+2\epsilon} G(p)^n dp \right]}{f_{L,\epsilon}(1-\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp} + \frac{F(\tilde{\gamma}-2\epsilon)G(\tilde{\gamma}-2\epsilon)^n + (1-F(\tilde{\gamma}+2\epsilon))G(\tilde{\gamma}+2\epsilon)}{f_{L,\epsilon}(1-\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp}$$

The right hand side of Equation (5) has five terms. These will be considered in turn. First:

$$\frac{f_{H,\epsilon}(\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp}{f_{L,\epsilon}(1-\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp} = \frac{f_{H,\epsilon}(\tilde{\gamma})}{f_{L,\epsilon}(1-\tilde{\gamma})}$$

Next consider denominator of the second term. $G(p)$ is increasing from $\tilde{\gamma} - \epsilon$ to $\tilde{\gamma}$ and hence:

$$\int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp \geq \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}} G(p)^n dp \geq \epsilon G(\tilde{\gamma}-\epsilon)^n$$

Taking the second term, and allowing $n \rightarrow \infty$, it follows that:

$$\frac{f_{H,2\epsilon}(\tilde{\gamma}) \int_{\tilde{\gamma}-2\epsilon}^{\tilde{\gamma}-\epsilon} G(p)^n dp}{\int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp} \leq \frac{f_{H,2\epsilon}(\tilde{\gamma})}{\epsilon} \int_{\tilde{\gamma}-2\epsilon}^{\tilde{\gamma}-\epsilon} \left[\frac{G(p)}{G(\tilde{\gamma}-\epsilon)} \right]^n dp \rightarrow 0$$

which holds since $G(p) < G(\tilde{\gamma} - \epsilon)$ for all $p < \tilde{\gamma} - 2\epsilon$. An identical argument ensures that the third term vanishes. A similar argument applies to the fourth term:

$$\frac{F(\tilde{\gamma}-2\epsilon) G(\tilde{\gamma}-2\epsilon)^n}{f_{L,\epsilon}(1-\tilde{\gamma}) \int_{\tilde{\gamma}-\epsilon}^{\tilde{\gamma}+\epsilon} G(p)^n dp} \leq \frac{F(\tilde{\gamma}-2\epsilon)}{\epsilon f_{L,\epsilon}(1-\tilde{\gamma})} \left[\frac{G(\tilde{\gamma}-2\epsilon)}{G(\tilde{\gamma}-\epsilon)} \right]^n \rightarrow 0$$

with a symmetric argument for the fifth term. Conclude from this that:

$$\lim_{n \rightarrow \infty} r(\tilde{\gamma}) \leq \frac{f_{H,\epsilon}(\tilde{\gamma})}{f_{L,\epsilon}(1-\tilde{\gamma})}$$

Notice now that ϵ may be chosen arbitrarily small. It follows that:

$$\lim_{n \rightarrow \infty} r(\tilde{\gamma}) \leq \lim_{\epsilon \rightarrow 0} \frac{f_{H,\epsilon}(\tilde{\gamma})}{f_{L,\epsilon}(1-\tilde{\gamma})} = \frac{f(\tilde{\gamma})}{f(1-\tilde{\gamma})}$$

A symmetric procedure bounds the limit below, and hence $r(\tilde{\gamma}) \rightarrow f(\tilde{\gamma})/f(1-\tilde{\gamma})$. Next, construct a compact interval $[\gamma - \psi, \gamma + \psi]$ around γ , for small ψ . For $\tilde{\gamma} \in [\gamma - \psi, \gamma + \psi]$, the argument above establishes that $r(\tilde{\gamma}) \rightarrow f(\tilde{\gamma})/f(1-\tilde{\gamma})$ pointwise on this interval. But since $r(\tilde{\gamma})$ and its limit are continuous, and the interval is compact, it follows that this convergence

is uniform. Now, recall that $\gamma_n = \lceil \gamma n \rceil / n$, and hence $\gamma_n \rightarrow \gamma$. It follows that $\gamma_n \in [\gamma - \psi, \gamma + \psi]$ for sufficiently large n . For sufficiently large n , $r(\gamma_n)$ is arbitrarily close to $f(\gamma_n)/f(1 - \gamma_n)$. But by taking ψ sufficiently small, it is assured that this is arbitrarily close to $f(\gamma)/f(1 - \gamma)$, which follows from the continuity of f . This completes the proof. \square

A.2. Affine Belief Rules. This section provides omitted proofs from Sections 5 and 6.

Proof of Lemma 3: Consider an arbitrary smoothly increasing belief rule $\lambda(\delta_i)$. Define:

$$p = \Pr[\lambda(\delta_i) + \eta + \epsilon_i \geq 0 \mid \eta] = H(\eta)$$

This is the probability that a randomly selected agent i supports option 1, given the common utility component η . Given that $\lambda(\delta_i)$ is smoothly increasing, $H(\eta)$ is strictly and smoothly increasing in η . Write $h(\eta) = H'(\eta)$. It follows that:

$$F(p) = \Pr[\eta \leq H^{-1}(p)] = \Phi \left(\frac{H^{-1}(p) - \mathbb{E}[\eta]}{\sqrt{\text{var}[\eta]}} \right)$$

where Φ is the cumulative distribution function of the normal. This probability (via $\mathbb{E}[\eta]$ and $\text{var}[\eta]$) is conditional on the information available to the agent $i = 0$, and uses the fact that posterior beliefs over η are normal (Lemma 1). Differentiate to obtain:

$$f(p) = \frac{1}{h(H^{-1}(p))\sqrt{\text{var}[\eta]}} \phi \left(\frac{H^{-1}(p) - \mathbb{E}[\eta]}{\sqrt{\text{var}[\eta]}} \right)$$

Next, using Proposition 1 and noting that Assumption 5 is satisfied:

$$\lim_{n \rightarrow \infty} \frac{q_1}{q_2} = \frac{f(\gamma)}{f(1 - \gamma)}$$

Evaluate this expression and take logs to obtain:

$$\begin{aligned} (6) \quad \log \frac{f(\gamma)}{f(1 - \gamma)} &= \log \frac{h(H^{-1}(1 - \gamma))}{h(H^{-1}(\gamma))} - \frac{(H^{-1}(\gamma) - \mathbb{E}[\eta])^2}{2\text{var}[\eta]} + \frac{(H^{-1}(1 - \gamma) - \mathbb{E}[\eta])^2}{2\text{var}[\eta]} \\ &= \log \frac{h(H^{-1}(1 - \gamma))}{h(H^{-1}(\gamma))} + \frac{H^{-1}(1 - \gamma)^2 - H^{-1}(\gamma)^2}{2\text{var}[\eta]} + \frac{H^{-1}(\gamma) - H^{-1}(1 - \gamma)}{\text{var}[\eta]} \mathbb{E}[\eta] \end{aligned}$$

Notice that this is affine in $\mathbb{E}[\eta]$. Now, voter $i = 0$ bases her beliefs on the privately observed signal δ_0 . Employ Lemma 1 from Section 2 to yield:

$$\eta \sim N \left(\frac{\kappa^2 \mu + \sigma^2 \delta_0}{\kappa^2 + \sigma^2}, \frac{\kappa^2 \sigma^2}{\kappa^2 + \sigma^2} \right)$$

Thus $\text{var}[\eta]$ does not depend on δ_0 , and $\mathbb{E}[\eta]$ is affine in δ_0 , so that $\hat{\lambda}(\delta_0)$ is affine in δ_0 . \square

Proof of Lemma 4: Suppose that each agent $i \geq 1$ adopts the posited affine belief rule, so that $\lambda_i = a + b\delta_i$. Individual i votes for option 1 whenever:

$$\eta + a + b\delta_i + \epsilon_i \geq 0 \quad \Leftrightarrow \quad a + (1 + b)\eta \geq -\epsilon_i - b(\delta_i - \eta)$$

Conditional on η , the right hand side is normally distributed with zero expectation, and variance $\tilde{\kappa}^2 = \text{var}[\epsilon_i + b(\delta_i - \eta)]$. It follows that:

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

where Φ is the cumulative distribution function of the normal. Differentiate to obtain:

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

Begin with the first term of Equation (6). First 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:

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

Next consider the second term of Equation (6).

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

Similarly:

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

It follows that:

$$\frac{H^{-1}(1-\gamma)^2 - H^{-1}(\gamma)^2}{2\text{var}[\eta]} = \frac{2a\tilde{\kappa}\Phi^{-1}(\gamma)}{\text{var}[\eta](1+b)^2}$$

The final term is simply:

$$\frac{H^{-1}(\gamma) - H^{-1}(1-\gamma)}{\text{var}[\eta]} \text{E}[\eta] = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)(1+b)}{(1+b)^2\text{var}[\eta]} \text{E}[\eta]$$

Assembling, obtain:

$$\log \frac{f(\gamma)}{f(1-\gamma)} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)(a + (1+b)\text{E}[\eta])}{\text{var}[\eta](1+b)^2}$$

The next step is the evaluation of the expectation and variance of η , conditional on the information the focal agent $i = 0$. Once again, recalling Lemma 1 of Section 2 it follows that:

$$\text{E}[\eta] = \frac{\kappa^2\mu + \sigma^2\delta_0}{\kappa^2 + \sigma^2} \quad \text{and} \quad \text{var}[\eta] = \frac{\kappa^2\sigma^2}{\kappa^2 + \sigma^2}$$

Substitute in to obtain:

$$\hat{\lambda}(\delta_0) = \log \frac{f(\gamma)}{f(1-\gamma)} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)(a(\kappa^2 + \sigma^2) + (1+b)(\kappa^2\mu + \sigma^2\delta_0))}{\kappa^2\sigma^2(1+b)^2}$$

Separate this out to obtain the affine function:

$$\hat{\lambda} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)[a(\kappa^2 + \sigma^2) + (1+b)\kappa^2\mu]}{\kappa^2\sigma^2(1+b)^2} + \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)\delta_0}{\kappa^2(1+b)}$$

Taking the intercept:

$$\begin{aligned} \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)[a(\kappa^2 + \sigma^2) + (1+b)\kappa^2\mu]}{\kappa^2\sigma^2(1+b)^2} &= \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)\sigma^2[a(\kappa^2 + \sigma^2) + (1+b)\kappa^2\mu]}{\kappa^2\sigma^2(1+b)\sigma^2(1+b)} \\ &= \frac{\hat{b}[a(\kappa^2 + \sigma^2) + (1+b)\kappa^2\mu]}{\sigma^2(1+b)} \end{aligned}$$

This yields the desired result. \square

Proof of Lemma 5: From Lemma 4 the best response $\hat{b}(b)$ satisfies:

$$\hat{b} = \frac{2\tilde{\kappa}\Phi^{-1}(\gamma)}{\kappa^2(1+b)}$$

This requires evaluation of $\tilde{\kappa}$, which satisfies:

$$\tilde{\kappa}^2 = \text{var}[\epsilon_i + b(\delta_i - \eta) \mid \eta] = \xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b$$

where ρ is the correlation coefficient between ϵ_i and δ_i , conditional on the common utility component η . Hence:

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

Evaluating the derivative:

$$\begin{aligned} \hat{b}'(b) &= \frac{2\Phi^{-1}(\gamma)}{\kappa^2} \left\{ \frac{b\kappa^2 + \rho\kappa\xi}{(1+b)\sqrt{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b}} - \frac{\sqrt{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b}}{(1+b)^2} \right\} \\ &= \frac{2\Phi^{-1}(\gamma)\sqrt{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b}}{\kappa^2(1+b)} \left\{ \frac{b\kappa^2 + \rho\kappa\xi}{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b} - \frac{1}{1+b} \right\} \\ &= \hat{b}(b) \left\{ \frac{b\kappa^2 + \rho\kappa\xi}{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b} - \frac{1}{1+b} \right\} \end{aligned}$$

This is decreasing for $b \geq 0$ if:

$$\frac{b\kappa^2 + \rho\kappa\xi}{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b} \leq \frac{1}{1+b}$$

Re-arrange this expression to obtain $\xi(\rho\kappa - \xi) \leq b\kappa(\rho\xi - \kappa)$. To check this inequality, first consider the right hand side. First $\rho \geq \kappa/\xi$ by assumption — see Section 2. Since $b \geq 0$, it is sufficient to show that the left hand side is weakly negative, which requires $\xi \geq \rho\kappa$. But this holds, since $0 \leq \rho \leq 1$ and $\kappa \leq \xi$. It follows that the function is (weakly) decreasing everywhere. Next, evaluate at the extremes to obtain:

$$\hat{b}(0) = \frac{2\xi\Phi^{-1}(\gamma)}{\kappa^2} \quad \text{and} \quad \lim_{b \rightarrow \infty} \hat{b}(b) = \frac{2\Phi^{-1}(\gamma)}{\kappa}$$

These calculations yield the desired properties of the function. \square

Proof of Lemma 6: To obtain an upper bound for the fixed point b^* , write $\hat{b}(b)$ as:

$$\hat{b}(b)\kappa = \frac{2\Phi^{-1}(\gamma_1)\sqrt{\xi^2 + b^2\kappa^2 + 2\rho\xi\kappa b}}{\kappa + b\kappa}$$

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

$$\hat{\beta}(\beta) = \frac{2\Phi^{-1}(\gamma_1)\sqrt{\xi^2 + \beta^2 + 2\rho\xi\beta}}{\kappa + \beta} \leq \frac{2\Phi^{-1}(\gamma)(\xi + \beta)}{\beta}$$

An upper bound may now be obtained by solving the equation:

$$\beta^2 - 2\Phi^{-1}(\gamma)(\xi + \beta) = 0$$

This has a positive root at:

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

It follows that an upper bound for the fixed point is:

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

This upper bound was obtained by setting $\rho = 1$. A tighter bound is available via a formal implementation of the microfoundation for the privately observed signal. In that case, the correlation coefficient satisfied $\rho = \kappa/\xi$. The bound on the transformed equation becomes:

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

It is clear that the right hand side is decreasing in κ . Hence sending $\kappa \downarrow 0$:

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

To obtain an upper bound, solve the equation:

$$\beta^4 - (2\Phi^{-1}(\gamma))^2(\xi^2 + \beta^2) = 0$$

This equation is quadratic in β^2 , and may be solved to obtain the positive root:

$$\begin{aligned} \beta^2 &= \frac{(2\Phi^{-1}(\gamma))^2 + \sqrt{(2\Phi^{-1}(\gamma))^4 + 4(2\Phi^{-1}(\gamma))^2\xi^2}}{2} \\ &= \frac{(2\Phi^{-1}(\gamma))^2}{2} \left\{ 1 + \sqrt{1 + \frac{\xi^2}{(\Phi^{-1}(\gamma))^2}} \right\} \end{aligned}$$

It follows that an upper bound is:

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

Moreover, it is clear that this bound is attained as $\kappa \downarrow 0$. □

Proof of Lemma 7: Consider the mapping:

$$B(b) = \hat{b}^{(2)}(b) = \hat{b}(\hat{b}(b))$$

Notice that \hat{b} is also a fixed point of B . Taking the derivative of this function:

$$B'(b) = \hat{b}'(\hat{b}(b))\hat{b}'(b)$$

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:

$$\begin{aligned} B'(b) &= b \left\{ \frac{\hat{b}\kappa^2 + \rho\kappa\xi}{\xi^2 + \hat{b}^2\kappa^2 + 2\rho\kappa\xi\hat{b}} - \frac{1}{1 + \hat{b}} \right\} \times \hat{b} \left\{ \frac{b\kappa^2 + \rho\kappa\xi}{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b} - \frac{1}{1 + b} \right\} \\ &= \left\{ \frac{\hat{b}^2\kappa^2 + \rho\kappa\xi\hat{b}}{\xi^2 + \hat{b}^2\kappa^2 + 2\rho\kappa\xi\hat{b}} - \frac{\hat{b}}{1 + \hat{b}} \right\} \times \left\{ \frac{b^2\kappa^2 + \rho\kappa\xi b}{\xi^2 + b^2\kappa^2 + 2\rho\kappa\xi b} - \frac{b}{1 + b} \right\} \end{aligned}$$

It is clear that, for $\rho > 0$, both of these terms are less than one, 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. \square

REFERENCES

- COX, G. W. (1994): "Strategic Voting Equilibria under the Single Nontransferable Vote," *American Political Science Review*, 88(3), 608–621.
- (1997): *Making Votes Count*. Cambridge University Press, Cambridge.
- DEGROOT, M. H. (1970): *Optimal Statistical Decisions*. McGraw-Hill, New York.
- DUVERGER, M. (1954): *Political Parties: Their Organization and Activity in the Modern State*. Wiley, New York.
- EVANS, G., J. CURTICE, AND P. NORRIS (1998): "New Labour, New Tactical Voting? The Causes and Consequences of Tactical Voting in the 1997 General Election," in *British Elections and Parties Review Volume 8: The 1997 General Election*, ed. by D. Denver, J. Fisher, P. Cowley, and C. Pattie, pp. 65–79. Frank Cass, London and Portland, OR.
- HOFFMAN, D. T. (1982): "A Model for Strategic Voting," *SIAM Journal of Applied Mathematics*, 42, 751–761.
- MCKELVEY, R., AND P. ORDESHOOK (1972): "A General Theory of the Calculus of Voting," in *Mathematical Applications in Political Science*, ed. by J. F. Herndon, and J. L. Bernd. University Press of Virginia, Charlottesville, volume 6.
- MORRIS, S., AND H. S. SHIN (1999): "Coordination Risk and the Price of Debt," mimeo Nuffield College, Oxford.
- MYATT, D. P. (1999): "A New Theory of Strategic Voting," Nuffield College Discussion Paper 1999-W21.
- MYERSON, R. B., AND R. J. WEBER (1993): "A Theory of Voting Equilibria," *American Political Science Review*, 87(1), 102–114.
- PALFREY, T. R. (1989): "A Mathematical Proof of Duverger's Law," in *Models of Strategic Choice in Politics*, ed. by P. C. Ordeshook, pp. 69–92. University of Michigan Press, Ann Arbor.
- SIMON, H. A. (1954): "Bandwagon and Underdog Effects and the Possibility of Equilibrium Predictions," *Public Opinion Quarterly*, 18, 245–253.