

Information Aggregation in Two-Candidate Elections

John O. Ledyard

It has been difficult to explain the relevance of public opinion polls, reporters, and interest-group endorsements in an election campaign—indeed it is difficult to understand the purpose of a campaign itself—within the context of standard spatial competition models. The problem is that these mechanisms have no role in the election process when every voter and candidate is fully informed from the beginning of the process. Only in the presence of imperfectly and asymmetrically informed agents can campaigns and their associated events have an impact on election processes and outcomes.¹

To date, modeling of elections with imperfectly informed players has generally² taken one of two approaches I refer to as the (1) naive and (2) rational expectations models. Neither, however, allows a complete and wholly consistent analysis of information processes in elections. In the naive approach, illustrated by Ledyard (1984) and Palfrey (in this volume), voters and candidates are modeled as Bayesians who use common knowledge about each other's priors and rationality to anticipate all actions and to compute their expected utility maximizing strategies. These are normal form models that sidestep issues that would arise in an extensive form. Although the equilibrium strategies can be complex contingent plans, such models are naive because they ignore an important possibility for the acquisition of information by uninformed agents. They assume that players do not look at the actual actions of the others and use that data, and the knowledge of the equilibrium strategy, to make better inferences about what the others know and, therefore, better inferences about the true state of the world. This assumption makes sense if all players move only once and simultaneously. However, as in most elections, if candidates can continually revise campaign strategies up to the actual vote and if voters can change intentions on the basis

I would like to thank John Ferejohn, who once asked me whether elections aggregated information like markets do. (The answer, John, is no.) I would also like to thank Richard McKelvey for helpful comments.

1. For an excellent survey of models of imperfect information in politics, as of 1985, see Calvert (1986).

2. Banks (1987a) and Harrington (1988b) are exceptions.

of endorsements and polls and other information about voter preferences and candidate plans, then it is necessary to recognize that a standard normal form model, which does not incorporate the possibility of sequential adjustment in the specification of strategies, is not sufficiently rich.

Recognizing the fact that voters and candidates can learn from each other's actual moves, but trying to preserve as much as possible from the complete information spatial competition model, McKelvey and Ordeshook (1985) take a second approach by adapting the rational expectations model of markets³ to elections. Both strategies and expectations are required to be in equilibrium simultaneously. That is, strategies are best replays given expectations and expectations are fully conditioned on the information revealed in the strategic choices.⁴ For example, if one candidate knows the median voter's ideal point then this candidate chooses the median as her platform. But the uninformed candidate will know, after observing that choice, exactly what the median position is and by choosing that median will ensure a tie with the informed candidate. The equilibrium is the same as that in the complete information model. This property of many rational expectations equilibria, that in equilibrium even initially uninformed players are fully informed, leads to the seeming paradox that there is no incentive to be informed because the other candidate will apparently be able to free ride on that information.⁵ The source of the paradox is the omission of an important strategic consideration in the equilibrium concept. In particular, there is no recognition of the fact that the informed candidate, knowing that the uninformed candidate will use the information in her strategy choice, may act to mask that information in an effort to mislead by choosing some position other than the median. Rational expectations equilibria may not remain as equilibria if a wider range of strategies is recognized.

In this essay we attempt to remedy the omission and to study elections in which candidates and voters can be imperfectly and privately informed, with a particular emphasis on the strategic effects of asymmetric information. We want to know whether information is aggregated,⁶ whether uninformed voters and candidates become informed, and to

3. In the economics literature on rational expectations in markets, players use prices to make inferences about others' information. See, e.g., Lucas (1972) or Plott and Sunder (1982).

4. For those not familiar with the rational expectations approach, the method by which information is revealed will be discussed in more detail below.

5. This paradox has also been noted in the economics analysis of markets and caused some problems for the efficient markets hypothesis.

6. We say information is aggregated when the behavior of agents, including their information inference activities, lead them to act together as if they are fully informed.

what extent information asymmetries cause outcomes to differ from those predicted by the models with completely informed players. As we will see below the answers depend on the institutional structure that is in place.

The essay is organized as follows. In sections I and II we introduce the model, while sections III–V contain the analysis. In section I, we summarize the structure needed from standard, complete information spatial competition models and give examples to illustrate the notation. In section II, we introduce the model of asymmetric information and define an equilibrium for the naive approach in which candidates can choose strategies conditional on their private information but not conditional on the strategy choice of the other candidate. In this simple context information is valuable because the better informed candidate will generally have a higher probability of winning the election. In section III, we look at what happens if candidates can make inferences about the other's private information from the other's strategic choices or, loosely, campaign behavior. We are especially interested in the sequence of strategy choices when there are no direct costs to changing positions, since when one candidate is fully informed and the other is not, the process is a signaling model. In section IV, we introduce a poll in order to generate public information in the midst of the strategic choices of the candidates. Again this is viewed as costless. In section V, we replace the poll with another election, so that there is a real opportunity cost to be paid if a candidate wants to mask his information. The current election can be lost trying to protect future chances. A summary of all the results and some observations are provided in section VI.

I. The Standard Election Model—Complete Information

To introduce notation and to provide a basis for the analysis we begin by summarizing the standard spatial competition model of two-candidate elections. Two candidates, A and B , are presumed to compete by selecting strategies in a space X , where X is a subset of the n -dimensional Euclidean space. These strategies are often called platforms but can include anything that affects voter behavior such as campaign strategies, speech topics and locations, expenditure decisions by state or category, timing of announcements, etc. Using a somewhat abusive notation we also let A be the strategy of candidate A and let B be the strategy of candidate B . Candidates are assumed to choose strategies to *maximize the probability of winning*.⁷ The aggregate behavior of individual voters

7. For these models, it can be shown that this is equivalent to the maximization of expected utility.

is captured in a function identifying the effect of candidate strategies on the probabilities of winning. Let $\text{prob}\{A \text{ wins}\} = \Pi(A, B, m) = 1 - \text{prob}\{B \text{ wins}\} = 1 - \Pi(B, A, m)$ where m parameterizes the voters' behavior.⁸ The zero-sum game, denoted by $G = \{A, B, X, \Pi, m\}$, is *the candidates' game*. An *equilibrium* of the game G is a pair of strategies (A', B') such that each is a best reply to the other. That is, $\Pi(A, B', m) \leq \Pi(A', B', m) \leq \Pi(A', B, m)$ for all $A, B \in X$. Note that $\Pi(A', B', m) = 1/2$.

To see how most specific spatial competition models fit into this framework let us examine three. In each, voters are assumed to have preferences over the strategy space captured in a utility function $U(x, e)$ where e are the parameters identifying each particular voter. One interpretation of U is that it is the indirect expected utility to the voter if the candidate using strategy x were to win. The entire collection of voters is then described by a density function $f(e, m)$, parameterized by m , where $f(e', m')$ describes the proportion of the voters whose preferences are defined by e' when the true parameter is m' . With this notation in mind we can turn to the examples and derive Π for each.

Example 1: The Median Voter (Downs 1957; Enelow and Hinich 1984b)

Strategies and the parameters e and m are real numbers where e is interpreted to be the ideal point of the voter and m is the median of the distribution of the ideal points. The utility functions are assumed to be single-peaked at e or, more specifically, $U(x, e) = -\|x - e\|$. It is also assumed that there are no abstentions. For this model of voter behavior, it is easy to show that

$$\Pi(A, B, m) = \begin{cases} 1 \\ 1/2 \\ 0 \end{cases} \text{ if } \|A - m\| \begin{cases} < \\ = \\ > \end{cases} \|B - m\|.$$

In equilibrium $A = B = m$. If the strategy space has more than 1 dimension an equilibrium rarely exists.

Example 2: Rational Abstentions (Ledyard 1984)

The strategy space is n -dimensional and the utility function of a typical voter is $U(x, e) = V(x, d) - c$ where $e = (d, c)$, c is the cost of voting for this voter, and d parameterizes the voters' preferences. If we assume that $f(e, m) = g(d, m)h(c)$ and if abstentions are rational (that is, a voter

8. The last equality assumes that the candidates' names are unimportant. The entire theory presented below can be easily generalized to the case in which names are important and the last equality does not hold.

abstains if and only if the expected benefits from voting do not outweigh the cost) then it can be shown that

$$\Pi(A, B, m) \begin{cases} > \frac{1}{2} & \text{if } \int V(A, d)g(d, m)dd > \int V(B, d)g(d, m)dd \\ = \frac{1}{2} & = \\ < \frac{1}{2} & < \end{cases}$$

In equilibrium $A = B = \operatorname{argmax}_x \int V(x, d)g(d, m)dd$. In equilibrium, no one votes but the threat of turnout still drives the candidates. A sufficient condition for equilibrium to exist is that the density of the costs of voting is uniform on $[0, C]$ where $C > 0$. Equilibrium also exists for other densities.

Example 3: Probabilistic Voting (Coughlin and Nitzan 1981)

The strategy space is n -dimensional. It is assumed that voters vote randomly but systematically and that $\operatorname{prob}\{e \text{ votes for } A\} / \operatorname{prob}\{e \text{ votes for } B\} = U(A, e) / U(B, e)$. There are no abstentions. For this model of voter behavior,

$$\Pi(A, B, e) = \int \operatorname{prob}\{e \text{ votes for } A\} f(e, m) d e.$$

In equilibrium $A = B = \operatorname{argmax}_x \int [\ln U(x, e)] f(e, m) d e$. Equilibrium exists for a wide range of densities.

II. The Standard Election Model—Incomplete Information

We want to study what happens if the candidates do not know m but instead have only some private and possibly indirect information about voters.⁹ The basis for the private information could be privately commissioned polls, local knowledge, prior beliefs, experience, etc. We use an approach that is standard in game and auction theory and model the election by assuming that candidates A and B privately observe signals, a and b . The triple (a, b, m) is assumed to be distributed according to the density function $g(a, b, m)$, which is common knowledge (as is Π). This is

9. We do not explicitly consider what the voters know but include their information in the function Π . Those interested in models with uninformed voters, and in the strategic revelation of information by candidates to voters, should consult Banks (1987a) and Harrington (1988b). Another model with uninformed voters can be found in McKelvey and Ordeshook (1985). They do not deal with strategic issues but instead postulate a rational expectations model of equilibrium behavior.

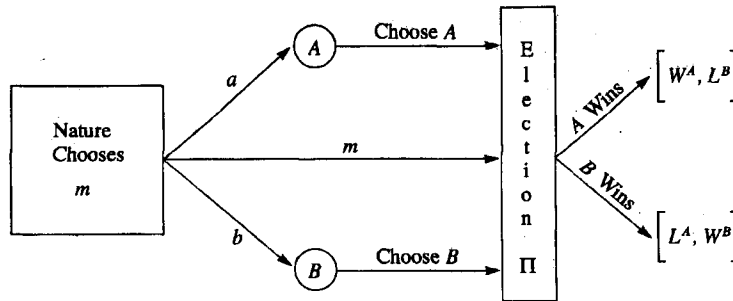


Fig. 1

a game with imperfect information in which nature moves first and picks a particular value of (a, b, m) . A and B , after seeing a and b respectively, simultaneously select their strategies, A and B . Voters then vote according to the function $\Pi(A, B, m)$ which yields the outcome—someone wins. An equilibrium is a strategy for each player describing what that candidate will do conditional on his or her private information (see fig. 1).

More formally, the (Bayes) game is $G = \{A, B, \Pi(A, B, m), g(a, b, m)\}$. A (contingent) strategy for candidate A is a function $A(a)$; for B it is a function $B(b)$. A (Bayes) equilibrium of G is a pair of functions $A^*(a), B^*(b)$ such that for every a

$$A^*(a) \in \operatorname{argmax}_A \int \Pi(A, B^*(b), m) g(m, a, b) dm db \tag{1.1}$$

and for every b

$$B^*(b) \in \operatorname{argmin}_B \int \Pi(A^*(a), B, m) g(m, a, b) dm da. \tag{1.2}$$

To develop a feel for this model let us look at three special cases: (1) symmetric information, (2) no information, and (3) extremely asymmetric information.

Example 4: Symmetric Information

To model a situation in which A and B are in symmetric situations, *ex ante*,¹⁰ with respect to information we assume that a and b are independently and identically distributed, *given* m . That is, we assume that $g(m, a, b) = g^1(a, m)g^1(b, m)$. For this informational structure there is a symmetric equilibrium in which $A^*(s) = B^*(s) = C^*(s)$ for every signal s . Further, the *ex ante* expected probability that A wins is

$$\int \Pi[C^*(a), C^*(b), m] g^1(a, m) g^1(b, m) da db dm.$$

10. *Ex ante* refers to the situation before a and b are known.

It is easy to see that this probability equals $\frac{1}{2}$ in the absence of any other asymmetries. It is also true, however, that the *interim*¹¹ expected probability that A wins given a , defined as

$$\int \Pi(A^*(a), B^*(b), m) g(a, b, m) db dm$$

can be anything between 0 and 1. A similar observation can be made with respect to the *ex post*¹² expected probability that A wins given a and b , which is

$$\int \Pi(A^*(a), B^*(b), m) g(a, b, m) dm.$$

Example 5: No Prior Information

To model a situation in which neither candidate possesses information we assume that a and b are distributed independently and identically and that both a and b are independent of m . In this case, $g(a, b, m) = r(a)r(b)t(m)$ and in equilibrium $A^*(a) = A^*$ and $B^*(b) = B^*$. That is, optimal strategies ignore the signals because, for example, given the signal a

$$A^*(a) \in \operatorname{argmax}_A \int \Pi(A, B^*(b), m) r(b) t(m) db dm$$

and the objective function is independent of a . One can simply redefine the probability of winning as $\Pi^*(A, B) = \int \Pi(A, B, m) t(m) dm$. The game is now as in section I with complete information. Neither player knows anything of importance that the other does not.

Example 6: Asymmetrically Informed Candidates

To illustrate the extent to which information about the electorate can be valuable for a candidate, consider a case in which A is well informed and B knows very little. This might occur because A is the incumbent and knows the district while B is the challenger with very little experience. To capture this asymmetry we let $g(a, b, m) = r(a, m)t(b)$. B 's signal provides no information while A 's signal tells A something about m . In effect, the marginal density on m captures what both A and B know about the distribution of voters while the correlation between a and m , described by $r(\cdot)$, determines the extent to which A is better informed than B . In equilibrium, $B^*(b) = \hat{B}$ and for every realization of a ,

11. *Interim* refers to the situation that exists when A knows a but not b or m and B knows b but not a or m .

12. *Ex post* describes the situation after a and b are known.

$\int \Pi [A^*(a), \hat{B}, m] r(m, a) dm \geq \frac{1}{2}$ (for most examples it is $> \frac{1}{2}$). In an extreme case, if A is fully informed, then

$$A^* = \hat{A} \text{ and } \hat{A} \in \operatorname{argmax}_A \Pi[A, \hat{B}, m].$$

Because

$$\hat{B} \notin \operatorname{argmax}_B \Pi[\hat{A}, B, m], \Pi[\hat{A}, \hat{B}, m] > \frac{1}{2},$$

and, therefore, information is valuable. This example illustrates that an incumbency advantage can be derived from the possession of better *private* information about the voting tendencies of the constituency to be served.

Since A 's equilibrium strategy depends on a , an outside observer could invert the relationship after observing the action A^* , compute $A^{*-1}(A^*) = \{a \mid A^*(a) = A^*\}$, and then use that information to update his prior. If $A^{*-1}(A^*) = a^*$, then the outside observer would know everything A does. If B could do this *before the election*¹³ then B could, presumably, minimize A 's informational advantage. The election process would be aggregating privately held information and the equilibrium might look like a rational expectations equilibrium. Counteracting this tendency is the fact that A , as a strategic player, will recognize that B can gain from information leakages and A will act to minimize the loss by choosing a different strategy.

There are at least two questions to be studied. Do elections aggregate information? Do informational considerations distort the strategic choices of candidates from those predicted by the standard model? As we will see, the answers to such questions depend on institutional features like the timing of strategic choices, the existence of public polls, and the existence of future elections.

III. The Timing of Strategic Choices

If the candidates can react to each other's electoral strategies they may be able to infer each other's private information which leads to information aggregation. This is the intuition behind rational expectations equilibria. To ascertain whether this intuition is valid in elections, we begin with the simplest election scenario. First candidates A and B simultaneously pick, respectively, strategies α and β . After observing each other's choices, they choose again by simultaneously picking, respectively, A and B , at which point the voters vote. The timing structure is shown in figure 2. The

13. We have assumed in (1.1) and (1.2) that A and B move simultaneously, so that any chance to infer A 's information arises too late to help B strategically.




Fig. 2

crucial assumption, which generates the main theorem in this section, is that the first moves, (α, β) , are free in the sense that voters vote only on the basis of A and B . In the economics literature this type of move has become known as *cheap talk* and its use can affect the outcome.¹⁴ To analyze this new game, we need a tighter equilibrium concept than Bayes equilibrium, which allows us to see what occurs when one agent can infer something about the other's private information. We use the idea of a sequential equilibrium, the main feature of which is a requirement of subgame perfection—a form of dynamic rationality.

More formally, the game is $G = \{A, B, \Pi(A, B, m)g(a, b, m)\}$. A (contingent) strategy for candidate A is a pair of functions $\alpha(a)$ and $A(\alpha, \beta, a)$, and a (contingent) strategy for B is a pair of functions $\beta(b)$ and $B(\alpha, \beta, b)$. A (sequential¹⁵ Bayes) equilibrium of G is the 4-tuple of functions $\alpha(a), \beta(b), A(\alpha, \beta, a), B(\alpha, \beta, b)$ such that for each a

$$\alpha(a) \in \underset{\alpha}{\operatorname{argmax}} \int \Pi[A(\alpha, \beta(b), a), B(\alpha, \beta(b), b), m]g(m, a, b)dmdb, \quad (2.1)$$

14. The usual example of a nonfree move is a change in platform so radical that voters switch to the other, more dependable, candidate. Banks (1987a) uses a somewhat *ad hoc* cost proportional to the distance $\|\alpha - A\|$ to model this. In our context an "opportunity cost" would arise if x were time indexed as $(x_1, x_2) = x$. Then the moves α and A would constitute parts of the strategy (α, A) . At time 2, α is fixed and can affect voters' reactions to A since $\Pi = \Pi[(\alpha, A), (\beta, B), m]$. We assume below in (2) that voters ignore the moves α, β , and let $\Pi = \Pi(A, B, m)$ to emphasize the signaling nature of the first moves. In section V, we allow the first moves to be costly.

15. Relations (2.1) to (2.4) are necessary but not sufficient for a sequential equilibrium. An additional criterion must be met that specifies how candidates will react to zero-probability events. That is, $A(\alpha, \beta, a)$ must also be defined on β for which there are no b such that $\beta = \beta(b)$. We do not do that here since the results we are interested in below do not depend in any crucial way on this extension. The reader is free to apply his or her favorite refinement.

for each b

$$\beta(b) \in \underset{\beta}{\operatorname{argmin}} \int \Pi[A(\alpha(a), \beta, a), B(\alpha(a), \beta, b), m]g(m, a, b)dm da, \quad (2.2)$$

for all α, a and for all β such that $\beta = \beta(b)$ for some b ,

$$A(\alpha, \beta, a) \in \underset{A}{\operatorname{argmax}} \int \Pi[A, B(\alpha, \beta, b), m]g(m, a, b | \beta = \beta(b))dm db, \quad (2.3)$$

and such that for all β, b and for all α such that $\alpha = \alpha(a)$ for some a ,

$$B(\alpha, \beta, b) \in \underset{B}{\operatorname{argmin}} \int \Pi[A(\alpha, \beta, a), B, m]g(m, a, b | \alpha = \alpha(a))dm da. \quad (2.4)$$

The key fact that we will observe in this model is that the equilibrium electoral outcomes are exactly the same as in the model of the previous section when no reaction is possible. That is, no information aggregation occurs and no advantage or disadvantage accrues to either candidate from her ability to react to the other. This is true since the equilibrium functions $\alpha(a)$ and $\beta(b)$ are constant (termed *pooling strategies*) and no inferences can be made about private information. The intuition is this: if, say, A were to choose α differently in different states of information, then B , knowing this and being able to condition on that information, would be able to use that information and do better. Since this is a zero-sum game, when B does better A does worse. Anything A says can and will be used against her. Of course, A can anticipate that B will use the information and can scramble the signals by permuting the strategic choice of α across signals. If B does not anticipate this then A will be better off because A would have misled B into using a suboptimal strategy for the true state. Of course, B prevents this by ignoring A 's α strategy. In equilibrium, A cannot gain by separating (choosing different α in different information sets) but can lose if B uses the information. In equilibrium, B cannot lose by ignoring any separation and can lose if A uses any attempt at inference by B against B .

For the reader who is unfamiliar with the machinery of Bayesian games, an example is helpful. Consider a simple form of Down's model (example 1 in section I) in which there are only two possible medians: $m \in \{L, R\}$. Assume that each candidate has only two strategies. That is, A, B, α , and β belong to the set $\{l, r\}$. (It is not necessary that either $l = L$ or $r = R$ but they can be interpreted as such.) Finally, assume that B knows the value of m and that A does not. A believes the probability

that $m = L$ is $\frac{1}{4}$. The payoff function can be summarized in a payoff table for A.

$$\Pi(A, B, L) = \begin{array}{c|cc} A \setminus B & l & r \\ \hline l & \frac{1}{2} & 1 \\ \hline r & 0 & \frac{1}{2} \end{array}$$

and

$$\Pi(A, B, R) = \begin{array}{c|cc} A \setminus B & l & r \\ \hline l & \frac{1}{2} & 0 \\ \hline r & 1 & \frac{1}{2} \end{array}$$

The two-move game can then be thought of as follows:

Move 0: Nature Chooses m

Move 1: B picks $\beta \in \{l, r\}$, A 's choice is unimportant since A is uninformed,

Move 2: A picks $A \in \{l, r\}$, B chooses $B = m$.

Those who are familiar with the literature will note that this is just a very simple signaling game whose game tree and information sets are given in figure 3. It is easy to compute that, even if mixed strategies are allowed (as long as A can differentiate between different mixtures), the only equilibria involve B pooling at move 1 and A picking r at move 2. B can pool by picking l whether the true state is L or R or by picking r whether the true state is L or R . Both strategies leave A uninformed. The equilibrium outcome is that B wins if $m = L$ and there is a tie if $m = R$. This is identical to the outcome of a one (simultaneous) move game. The ex ante probability that B wins, calculated before B receives information, is $\frac{3}{8}$ and the ex ante probability that A wins is $\frac{3}{8}$. One can interpret this as the advantage from being informed. It also indicates that there are gains from the acquisition of private information and suggests that an interesting variation, not studied in this essay, would be to let candidates buy information (e.g., pay for private polls) rather than simply begin with their information in hand.

To capture the simple intuition from this example in a more general theorem, we first need to handle some minor technical details that arise from the possibility of multiple equilibria. We will show that pooling at the first stage is an equilibrium strategy. That is, there are constant functions $\alpha(a) = \alpha'$ and $\beta(b) = \beta'$ which are the first part of a sequential equilibrium. There may also be other equilibria, but both candidates will




Fig. 3

be indifferent between these and the pooling equilibrium. To establish this, we must identify when nonpooling is inessential.

DEFINITION. Let $\alpha^*, \beta^*, A^*, B^*$ be an equilibrium satisfying (2). We say that α^* is an essentially pooled strategy (a symmetric definition applies for β^*) if and only if, for every $a, b, \alpha', \beta', \alpha'', \beta''$ such that $\alpha' = \alpha^*(a), \beta' = \beta^*(b)$, and $\alpha'' = \alpha^*(a'')$ for some a'' ,

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta', a), B^*(\alpha', \beta', b), m] g(m, a, b \mid \beta' = \beta^*(b)) dm db \\ &= \int \Pi[A^*(\alpha'', \beta', a), B^*(\alpha', \beta', b), m] g(m, a, b \mid \beta' = \beta^*(b)) dm db \end{aligned} \tag{3.1}$$

and

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta', a), B^*(\alpha', \beta', b), m] g(m, a, b \mid \alpha' = \alpha^*(b)) dm da \\ &= \int \Pi[A^*(\alpha', \beta', a), B^*(\alpha'', \beta', b), m] g(m, a, b \mid \alpha' = \alpha^*(a)) dm da. \end{aligned} \tag{3.2}$$

That is, both A and B are indifferent between using the second-move strategy generated by α' and that generated by α'' . Also both candidates are indifferent between the original equilibrium and that created by replacing α' with α'' .

We can now state the principal result of this section.

THEOREM 1. *If $\langle \alpha(a), \beta(b), A(\alpha, \beta, a), B(\alpha, \beta, b) \rangle$ is an equilibrium of the election game satisfying (2), then $\alpha(\cdot)$ and $\beta(\cdot)$ are essentially pooled strategies.*

Proof. Let $\alpha^*, \beta^*, A^*, B^*$ satisfy (2). We prove α^* is essentially pooled. The proof for β^* follows a symmetric argument. Let $Z = \{\alpha \mid \alpha^*(a) = \alpha \text{ for some } A\}$. If Z is a singleton then $\alpha^*(\cdot)$ is a pooled strategy and we are done. So let $\alpha', \alpha'', \epsilon Z$ with $\alpha' \neq \alpha''$, let $\alpha^*(a') = \alpha'$ and let $\alpha^*(a'') = \alpha''$. From (2.4) it follows that

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta', a), B^*(\alpha', \beta', b), m] g(m, a, b \mid \alpha' = \alpha^*(a)) dmda \\ & \leq \int \Pi[A^*(\alpha', \beta', a), B^*(\alpha'', \beta', b), m] g(m, a, b \mid \alpha' = \alpha^*(a)) dmda \end{aligned} \quad (4.1)$$

for all $\langle \beta', b \rangle$ such that $\beta' = \beta^*(b)$. If we integrate (4.1) over all b we get (remembering that $\beta' = \beta^*(b)$),

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta^*(b), a), B^*(\alpha', \beta^*(b), m)] g(m, a, b \mid \alpha' = \alpha^*(a)) dmdadb \\ & \leq \int \Pi[A^*(\alpha', \beta^*(b), a), B^*(\alpha'', \beta^*(b), m)] \times \\ & \quad g(m, a, b \mid \alpha' = \alpha^*(a)) dmdadb. \end{aligned} \quad (4.2)$$

Now (2.3) implies that, for all β' such that $\beta' = \beta^*(b)$ for some b ,

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta', \bar{a}), B^*(\alpha'', \beta', b), m] g(m, \bar{a}, b \mid \beta' = \beta^*(b)) dmdb \\ & \leq \int \Pi[A^*(\alpha'', \beta', \bar{a}), B^*(\alpha'', \beta', b), m] \times \\ & \quad g(m, \bar{a}, b \mid \beta' = \beta^*(b)) dmdb. \end{aligned} \quad (4.3)$$

Integrating (4.3) over all b implies

$$\begin{aligned} & \int \Pi[A^*(\alpha', \beta^*(b), \bar{a}), B^*(\alpha'', \beta^*(b), b), m] g(m, \bar{a}, b) dmdb \\ & \leq \int \Pi[A^*(\alpha'', \beta^*(b), \bar{a}), B^*(\alpha'', \beta^*(b), b), m] g(m, \bar{a}, b) dmdb. \end{aligned} \quad (4.4)$$

Integrate (4.4) over $\{a \mid \alpha^*(a) = \alpha'\}$ and use (4.2) to get

$$\begin{aligned} & \int \Pi^*[A^*(\alpha', \beta^*(b), a), B^*(\alpha', \beta^*(b), b), m] \times \\ & \quad g(m, a, b \mid \alpha' = \alpha^*(a)) dmdbda \\ & \leq \int \Pi[A^*(\alpha'', \beta^*(b), a), B^*(\alpha'', \beta^*(b), b), m] \times \\ & \quad g(m, a, b \mid \alpha' = \alpha^*(a)) dmdbda. \end{aligned} \quad (4.5)$$

Now we must consider two cases.

Case 1. $\exists b', \beta', \bar{a}', \alpha'', \alpha', \alpha''$ such that $\alpha^*(a') = \alpha' \neq \alpha^*(a'') = \alpha''$ and

$\beta' = \beta^*(b')$ and at least one of the inequalities (4.1) or (4.3) is strict. Then (4.5) is strict and \exists at least one a^* such that $\alpha' = \alpha^*(a^*)$ and

$$\alpha' \notin \underset{\alpha}{\operatorname{argmax}} \int \Pi[A^*(\alpha, \beta^*(b), a^*), B^*(\alpha, \beta^*(b), b), m] g(m, a, b) dm db.$$

Therefore, (2.1) is contradicted which proves the theorem.

Case 2. For all $a', a'', \alpha', \alpha'', \beta', b'$ such that $\alpha' = \alpha^*(a'), \beta' = \beta^*(b'), \alpha'' = \alpha^*(a'')$ the inequalities (4.1) and (4.3) are equalities. Then by (3.1) and (3.2), $\alpha^*(\cdot)$ and $\beta^*(\cdot)$ are essentially pooled strategies.

(Q.E.D.)

Although the ability to react to the strategy of one's electoral opponent and to take advantage of inferences about the opponent's private information is *potentially* beneficial, in equilibrium no leakage of information occurs because a candidate can only lose by allowing that to happen. The outcome of the electoral process is the same as it was in the one-move model of section II. No aggregation of information occurs. It can also be shown that increasing the number of allowed responses or moves in the election process to more than two will not change the conclusions. The only strategy choice that will depend, in equilibrium, on one's private information is that last one. Nothing important happens until the last move. Rational expectations are *not* an equilibrium.

COROLLARY 1.1. *The outcome of a T-move election (for T finite) is the same as the outcome of a one-move election.*

Finally, it is important to observe that the simultaneity of the moves is crucial to some of the conclusions. Returning to the simple signaling example with two medians, if B had to choose B first and A could choose A after observing B 's choice then there are two equilibria.

Equilibrium 1. B pools and picks r , A reacts by picking r also. In this case, the probability that A wins is $\frac{1}{2}$ for each value of m . B 's informational advantage has been lost because of the first move.

Equilibrium 2. B separates and chooses m . A matches B . Again the probability that A wins is $\frac{1}{2}$ for each value of m and B has lost the informational advantage.

Candidate A can guarantee at least a tie by matching B 's move and, therefore, B neither gains nor loses from the revelation of the information. Whether information is aggregated in this case is indeterminate.

IV. Polls

A very natural question to ask is whether the existence of public polls effect the aggregation of information. We modify the model in section III by having a poll taken and the results publicly announced between the candidates' first and second moves. First, nature determines the




Fig. 4

private information by picking a and b . Next A and B simultaneously choose and announce α and β . Then a poll is taken and the result p , which is the proportion of the voters who would vote for A , is announced. While it would not be difficult to allow for polling errors, we avoid the extra notation and assume that the poll result is governed by a nonrandom relationship $p = \rho(\alpha, \beta, m)$ which is common knowledge. After the poll results are announced, A and B choose strategies A and B , respectively. In accordance with the timing sequence shown in figure 4, the election is then held and the outcome occurs.

A (sequential) equilibrium is a 4-tuple of conditional strategies $\alpha(a), \beta(b), A(\alpha, \beta, p, a), B(\alpha, \beta, p, b)$ where

$$A(\cdot) \in \operatorname{argmax}_A \int \Pi(A, B(\cdot), m) \times g(m, a, b \mid \beta = \beta(b), p = \rho(\alpha, \beta, m)) dm db \quad (5.1)$$

$$B(\cdot) \in \operatorname{argmin}_B \int \Pi(A(\cdot), B, m) \times g(m, a, b \mid \alpha = \alpha(a), p = \rho(\alpha, \beta, m)) dm da \quad (5.2)$$

$$\alpha(a) \in \operatorname{argmax}_\alpha \int \Pi[A(\alpha, \beta(b), \rho(\alpha, \beta(b), m), a), B(\alpha, \beta(b), \rho(\alpha, \beta(b), m), b), m)] g(m, a, b) dm db \quad (5.3)$$

$$\beta(b) \in \operatorname{argmin}_\beta \int \Pi[A(\alpha(a), \beta, \rho(\alpha(a), \beta, m), a), B(\alpha(a), \beta, \rho(\alpha(a), \beta, m), b), m)] g(m, a, b) dm da. \quad (5.4)$$

The only difference between these relations and those in (2) is the fact that the choice of A and B is now conditioned on the information in the poll result, p , which can be affected by the choices of α and β .

Since the equilibrium conditions (5) are notationally complex and since the result we want to highlight is easier to understand in an example, we turn to a special case of the median voter model which is a bit more complex than that of section III in that we allow three possible

medians, $m \in \{L', M', R'\}$. Each candidate has three choices available where A, B, α , and β all belong to the set $\{L, M, R\}$. (It is not necessary that either $L' = L, M' = M$, or $R' = R$ but our definition of Π will use that identification.) Finally, assume A knows the value of m and B does not but it is common knowledge that B believes the probability that $m = L'$ is $\frac{1}{3}$, the probability that $m = M'$ is $\frac{1}{2}$, and the probability that $m = R'$ is $\frac{1}{6}$. For our example,

$$\Pi(A, B, m) =$$

AB	if $m = L'$			if $m = M'$			if $m = R'$				
	L	M	R	L	M	R	L	M	R		
L	$\frac{1}{2}$	1	1	L	$\frac{1}{2}$	0	$\frac{1}{2}$	L	$\frac{1}{2}$	0	0
M	0	$\frac{1}{2}$	1	M	1	$\frac{1}{2}$	1	M	1	$\frac{1}{2}$	0
R	0	0	$\frac{1}{2}$	R	$\frac{1}{2}$	0	$\frac{1}{2}$	R	1	1	$\frac{1}{2}$

Finally, let the function that determines the poll result be the same as the function that determines the electoral outcome. That is, let¹⁶ $p = \rho(\alpha, \beta, m) \equiv \Pi(\alpha, \beta, m)$.

To provide a benchmark, let us first examine the outcome when there is no poll. Using the model in section III, it is easy to compute that the equilibrium would involve A pooling in the choice of α . At the second move, A chooses L if $m = L'$, M if $m = M'$, and R if $m = R'$. B 's choice of β is unimportant (B is uninformed). B chooses M at the second move. The ex ante probability that A , the informed candidate, wins is $\frac{3}{4}$.

To understand the equilibrium with the poll, we need to work our way back from the last move to the first. At the second move, A chooses m no matter what has happened earlier. B extracts as much information as possible from the signals b, α, β , and p and then chooses the position with the highest probability of being the true median. More formally, for each possible pair of first period choices, α and β , we can calculate B 's information after the poll and, therefore, B 's strategy choice in period 2. B 's information and strategy choices are summarized thus:

16. In the Downs model and the Coughlin-Nitzan model this identification makes some sense if voters correctly reveal their intention to vote. In the Ledyard model of voter behavior it would not make sense since voting in the poll (without costs of participation) would follow the median voter model while voting in the election would not.

$\alpha \backslash \beta$	<i>L</i>	<i>M</i>	<i>R</i>
		L	L M
L	LMR	MR	R
	L		LM R
M	MR	LMR	R
	L M	LM	
R	R	R	LMR

where, for example, the entry for $(\alpha, \beta) = [L, M]$ means that if $m = L$ then B knows that $m = L$ and B plays L , but if $m = M$ or R then B knows that $m \in \{M, R\}$ and B plays M . (The horizontal configurations are the information sets and the bold-faced letter is the strategic choice.)

Given that we know A 's and B 's second move choices conditional on the true state m and the first move choices, we can now compute the best first move choices for each. The following table summarizes the best responses by A to a choice of β for each possible median.

median $\backslash \beta$	<i>L</i>	<i>M</i>	<i>R</i>
L'	L	MR	RM
M'	LMR	LMR	LMR
R'	LM	LM	R

Thus, if A knows that the median is R' and B has chosen β equal to L then A is indifferent between responding with an α of L or M . Since A is "better off" if B can learn nothing from A 's choice of α , A should pool. A quick examination of the previous table will show that A can pool with a best response by matching B 's choice. For example, if β is R then α should be R , no matter what the true median is. This matching and pooling choice also minimizes the information that B can extract from the process since $p = 1/2$ for every m when the first moves are the same. B learns nothing from the first moves or from the poll.

We can also calculate B 's best response to any α given that A pools. The following table summarizes the payoffs to A for each pair of first move choices.

$\alpha \backslash \beta$	<i>L</i>	<i>M</i>	<i>R</i>
<i>L</i>	$3/4^*$	$11/16$	$8/16^\circ$
<i>M</i>	$11/16$	$12/16^*$	$9/16^\circ$
<i>R</i>	$8/16^\circ$	$9/16$	$3/4^*$

We have identified A 's best response in each column with an x and have identified B 's best response in each row with an o . The important observation is that *there is no pure strategy equilibrium!* A always moves to prevent B from learning anything by matching B 's position and B always moves to separate his position from A 's to generate information. There is, however, a *mixed strategy equilibrium* to this game. In particular one can show, using the previous table, that if the probability that A chooses α equal to L is $\frac{1}{2}$ and equal to R is $\frac{1}{2}$, and if B uses the same mixed strategy, then that is an equilibrium. A still pools but B can extract, with positive probability, some information by using a mixed strategy. When α and β are different (which will happen $\frac{1}{2}$ the time) B will be able to infer the true median. The ex ante probability that A wins is $\frac{5}{8}$ which is less than the probability without the poll. A still retains an advantage from being better informed and A can still mask her private information by using a pooling strategy but her advantage is reduced through the existence of the public poll. An interesting study, which we do not pursue here, would analyze the tradeoffs between private and public polls, and the incentives for the provision of either, especially if polling is costly.

In the more general model the main features of the example survive. That is, as long as there are some first period moves by one candidate given the first move of the other candidate, which lead the poll to generate useful information for some m , then (1) there are no pure strategy equilibria, (2) candidates mask their own private information up to the election, but (3) information aggregation does occur. Let

$$\psi(\alpha', \beta', a) = \int \Pi[A(\alpha', \beta', \rho(\alpha', \beta', m), a), B(\alpha', \beta', \rho(\alpha', \beta', m), b), m] g(m, a, b) dm db$$

where $A(\cdot)$ and $B(\cdot)$ satisfy (5.1) and (5.2). Hence ψ is the indirect utility to A of the last period game given a . If there are multiple equilibria in (5), each candidate is indifferent between them because of the zero-sum nature of the game. Therefore, ψ is well defined.

DEFINITION: We say that the poll can generate useful information for A if for all β there is an α such that

$$\int \psi(\alpha, \beta, a) da > \int \psi(\beta, \beta, a) da.$$

For example, if ρ were constant for all α and β then it could generate no useful information. If there is some β , such that (6) is false for all α then B can, by choosing that β , ensure that the probability that he wins is at

least as high as it would be without the poll. A poll such that ρ is identical to Π , as in the example, can generate useful information for one of A or B as long as both are not already fully informed.

THEOREM. *If $\rho(\alpha, \alpha, m)$ is constant in m and if ρ can generate useful information then there is no equilibrium where α and β are pooled pure strategies.*

Proof. Suppose that α' and β' are part of an equilibrium. Then

$$\int \psi(\alpha', \beta', a) da \geq \int \psi(\beta', \beta', a) da = \int \phi(a) da,$$

where

$$\phi(a) = \int \Pi[A(a), B(b), m] g(m, a, b) dm db$$

and where $A(\cdot)$ and $B(\cdot)$ satisfy (1). Also, using a symmetric construction for B ,

$$\int \psi(\alpha', \beta', a) da \leq \int \psi(\alpha', \alpha', a) da = \int \phi(a) da.$$

Therefore

$$\int \psi(\alpha', \beta', a) da = \int \phi(a) da.$$

But if $\rho(\cdot)$ can generate useful information, there is an $\hat{\alpha}$ such that

$$\int \psi(\hat{\alpha}, \beta', a) da > \int \psi(\beta', \beta', a) da = \int \phi(a) da$$

Therefore there is $\hat{\alpha}$ such that

$$\int \psi(\alpha', \beta', a) da \geq \int \psi(\hat{\alpha}, \beta', a) da > \int \psi(\beta', \beta', a) da,$$

which contradicts (7).

(Q.E.D.)

COROLLARY. *Under the hypothesis of the theorem there is an equilibrium where α and β are mixed strategies and, for some realizations of those strategies, information is aggregated.*

With public polling, candidates that are not initially fully informed make their second period moves with more information than in the one-shot election analyzed in section II. An interesting question is whether a sequence of free moves, say $\alpha_1, \alpha_2, \dots$, and a sequence of polls, ρ_1, ρ_2, \dots , converge to full information revelation and, therefore, to a

fully informed rational expectations equilibrium. This remains an open question.

V. Sequences of Elections

To this point, all preliminary moves of the candidates are free of commitment or cost. That fact plus the zero-sum nature of the game leads candidates to choose pooling strategies up to the last move so as not to leak information to an opponent. If, on the other hand, an election is held after each move, opportunity costs become important. A candidate must assess the trade-off between the increase in the probability of winning the current election, by being responsive to the private information, and the decrease in the probability of winning the next elections, by revealing the information to the opponent. In contrast to our previous models, the existence of the trade-off can lead a candidate to a nonpooling strategy that allows some private information to pass to the opponent. The trade-off may also lead a candidate to not exploit fully an informational advantage today so as to preserve that advantage for tomorrow.

To see how nonpooling strategies are sustained, we use the simplest modification of the earlier models that allows us to study the phenomenon. The election scenario, described in figure 5, has A and B first choosing positions α and β , respectively. An election is then held and the outcome occurs according to the function $\Pi(\alpha, \beta, m)$. The outcome is common knowledge. The candidates then choose A and B , respectively, and another election determines the outcome according to $\Pi(A, B, m)$. We assume that m does not change between elections, although it is easy to accommodate a situation in which the m in period 1 and the m in period 2 are believed to be correlated.

Since sequences of elections have not been regularly analyzed in the literature, we must make some assumptions to complete the model. In particular we assume that (1) candidates have no time preference, (2) the value of winning is constant over time, and (3) the value of losing is constant over time. Taken together these imply that each expected utility-maximizing candidate wishes to maximize the sum of the probabilities of winning.

The equilibrium of the game is a 4-tuple of strategies $\langle \alpha(a), \beta(b), A(\alpha, \beta, p, a), B(\alpha, \beta, p, b) \rangle$ where p is the outcome of the first election, such that for all p such that $p = \Pi(\alpha, \beta, m)$ for some $\alpha = \alpha(a)$ and $\beta = \beta(b)$, it is true for all α, a , for all β such that $\beta = \beta(b)$ for some b

$$A(\cdot) \in \operatorname{argmax}_A \int \Pi[A, B(\alpha, \beta, p, b), m] \times$$

$$g(m, a, b \mid \beta = \beta(b), p = \Pi(\alpha, \beta, m)) dm db$$

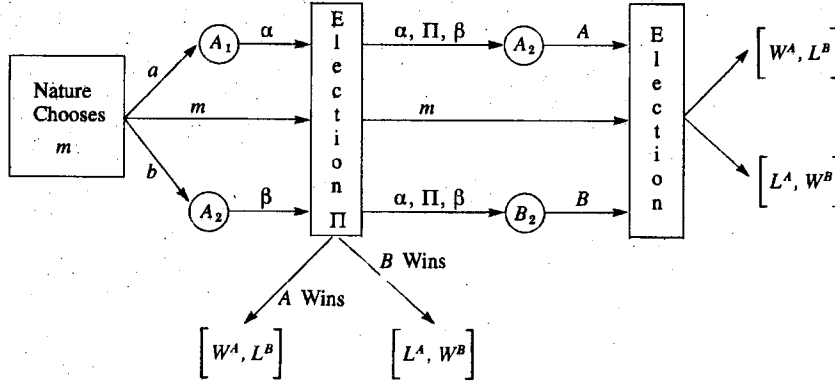


Fig. 5

for all β, b , for all α such that $\alpha = \alpha(a)$ for some a ,

$$B(\cdot) \in \operatorname{argmax}_B \int \Pi[A(\alpha, \beta, p, a), B, m] \times \\ g(m, a, b \mid \alpha = \alpha(a), p = \Pi(\alpha, \beta, m)) dmda$$

for all a

$$\alpha(\cdot) \in \operatorname{argmax}_\alpha \int \Pi(\alpha, \beta(b), m) + \Pi[A(\alpha, \beta(b), \Pi(\alpha, \beta(b), m), a), \\ B(\alpha, \beta(b), \Pi(\beta(b), m), b), m)] g(m, a, b) dmdb$$

for all b

$$\beta(\cdot) \in \operatorname{argmax}_\beta \int \Pi(\alpha(a), \beta, m) + \Pi[A(\alpha(a), \beta, \Pi(\alpha(a), \beta, m), m), \\ B(\alpha(a), \beta, \Pi(\alpha(a), \beta, m), m)] g(m, a, b) dmda.$$

To keep the analysis as simple as possible, we reexamine the example used in section III to see what new behavior can occur. Recalling that candidate A was fully informed as to which of three possible medians was the true one, it is easy to see that in the second election A will still fully exploit that information and will choose $A = L$ if $m = L'$, etc. The question is what will A do in the first election?

Suppose A follows a pooling strategy in equilibrium. We can compute A 's total payoff from each pair of first election strategies, α and β , where the computations of the choices A, B and the payoffs for the second election are identical to those in section IV. The total payoffs to A , conditional on α, β , and m , are

$\alpha \backslash \beta$	if $m = L'$			if $m = M'$			if $m = R'$		
	L	M	R	L	M	R	L	M	R
L	$\frac{1}{2} + 1^x$	$1 + \frac{1}{2}^x$	$1 + \frac{1}{2}$	$\frac{1}{2} + \frac{1}{2}$	$0 + \frac{1}{2}$	$\frac{1}{2} + \frac{1}{2}$	$\frac{1}{2} + 1$	$0 + 1$	$0 + \frac{1}{2}$
M	$0 + \frac{1}{2}$	$\frac{1}{2} + 1^x$	$1 + 1^x$	$1 + \frac{1}{2}^x$	$\frac{1}{2} + \frac{1}{2}^x$	$1 + \frac{1}{2}^x$	$1 + 1^x$	$\frac{1}{2} + 1^x$	$0 + \frac{1}{2}$
R	$0 + \frac{1}{2}$	$0 + 1$	$\frac{1}{2} + 1$	$\frac{1}{2} + \frac{1}{2}$	$0 + \frac{1}{2}$	$\frac{1}{2} + \frac{1}{2}$	$1 + \frac{1}{2}$	$1 + \frac{1}{2}^x$	$\frac{1}{2} + 1^x$

An x identifies A 's best replay(s) in each column. It is easy to see that the only possible pooling equilibrium is $\alpha = M$. But then, B 's best response is R . Therefore, pure pooling is not an equilibrium of the sequential election game. There is one partially pooling equilibrium: $\beta = M, \alpha = (M, M, R)$, and there is one separating equilibrium: $\beta = M, \alpha = (L, M, R)$. In this example, then, the private information held by A is simply not valuable enough to protect and so A chooses α to maximize the probability of winning election number 1. This tendency to separate would be reinforced if candidate A had a time preference (a discount rate less than 1) since then $\beta = M, \alpha = (L, M, R)$ is the only equilibrium.

It is not generally true, however, that candidates necessarily follow separating strategies in the first election. To see why, consider our example but with different prior beliefs. In particular, suppose it is common knowledge that B believes the prob ($m = L'$) = $\frac{1}{2}$, the prob ($m = M'$) = $\frac{3}{8}$, and the prob ($m = R'$) = $\frac{1}{8}$. Now the unique sequential equilibrium has first election components of $\beta = L$ and $\alpha = (L, M, M)$. This is a partially separating strategy since $\alpha(R) = M$. If A were to separate, then A would win the first election (if $m = M'$ or R') but A would reveal the information and tie the second. If A pools, then A wins the first when $m = M'$ or R' and also wins the second if $m = R'$ since B would choose M in the second.

In general, with many possible medians (say the whole real line) A can, by choosing just to the side of B on which the real median lies, protect her information *and* win the current election. Whether this can go on forever is unknown.

It does seem to be a general result that the informed candidate will not completely pool since full separation would be preferred to full-pooling. The intuition is simple for the case of a fully informed candidate versus an uninformed candidate. In a pooling equilibrium (since A must act as uninformed) the probability A wins in period 1 is $\frac{1}{2}$. Further, since some information may leak, the probability that A wins in period 2 is less than what it would be if no information leaked. If A were to fully separate in equilibrium then the probability A wins the first election is the same as if no information has leaked. The probability A wins the second is $\frac{1}{2}$ since both candidates are fully informed. A is therefore better off separating. Complete pooling is not an equilibrium.

A sequence of elections, then, can lead to the aggregation of information both from the election results and from a leakage of information from candidates who cannot afford to risk losing today's election in order to protect that information for the next election. A theorem remains to be provided for the general case.

VI. Summary and Observations

We have learned several facts about information in elections. First, and perhaps least surprising, we have seen that private information is valuable to candidates because it can raise the probability of being elected, in a symmetric equilibrium, to above .5.

Second, we found that, even if candidates can react to their opponent's strategic choices and make inferences about their opponent's private information, no information aggregation occurs because no candidate will follow a strategy that will allow her information to leak. They all follow an essentially pooling strategy. The positive benefit due the private information and the constant-sum nature of the election game provide strong incentives to each candidate to conceal whatever they know from the other. Even when a large, finite number of moves are available, the outcome remains identical to that which occurs when the candidates can move once simultaneously and, therefore, cannot react to their opponent's choices. Rational expectations equilibria will not occur. We also noted that if one candidate must move before the other then that candidate loses any informational advantage. If the informed candidate must move first then rational expectations equilibria may arise.

Third, we found that adding a public poll to the election process affects strategic choices and outcomes. Although each candidate continues to be able to prevent leakage of her private information through her (pooling) strategic choice, she cannot prevent her opponent from learning through the poll. One might think that the informed candidate could match exactly the strategy of her opponent and cause the poll to predict a tie and, therefore, to provide no useful information, but her opponent can defeat this tactic with a mixed strategy. Since the poll provides both candidates with the same information and since candidates can successfully hide their private information, the election outcome with a poll is the same as that of the single choice model in section II if we update each candidate's prior with the poll.

Fourth, however, the candidates will not be able to hide private information by using a pooling strategy in a sequence of elections. Doing so may lead the candidate to lose the current election; so there is a cost

to hiding information as well as a cost to revealing it. Although we prove no general results, we do provide an example that suggests two phenomena: (1) Some information aggregation will occur since it will be too costly to fully pool, and (2) candidates will skew their behavior in the sense that they will follow different strategies than those which would be best in a one-move model. A sequence of two elections with two privately informed candidates is not equivalent to two single elections and rational expectations may not occur.

Necessarily, we have several questions unanswered, and these questions are of three types. First, what occurs if the acquisition of the candidates' private information and the public polling information is costly? Because information aggregation is not complete, contrary to the assumption of a rational expectations equilibrium, there remains a benefit to its acquisition. What happens if the collection of private information is endogenous? If, in a simple model we allow private polls, in which reliability increases with the amount of money spent on it, we could compare the benefits from public polls to those of the private polls.

Second, what occurs if the moves of the candidates are costly? For example, we assume throughout that if a candidate uses a sequence of moves A_1, A_2, \dots, A_T , then the probability that the candidate wins is $\Pi(A_T, B_T, m)$ which does not depend on the early moves. That is, all early moves are "cheap talk." An alternate assumption is to suppose that the probability of winning is $\Pi(A_1, \dots, A_T, B_1, \dots, B_T, m)$. While partial answers are given by the example of sequential elections in section V and by Bernhardt-Ingberman (1985), Banks (1987a), and Harrington (1988b), a more thorough and less ad hoc analysis is required. This would allow us, for example, to better understand the informational benefits from incumbency as well as the costs of being "locked" into positions. We also need to explicitly merge models of uninformed candidates with models of uninformed voters, which requires a much better understanding of the process by which voters acquire information.

Third, what occurs if the number of moves in each sequence is increased? For example, we might conjecture that as the number of polls prior to the election is increased, the election outcome approaches the outcome predicted by a rational expectations equilibrium. A similar result might be available for a large sequence of elections but, since candidates distort their strategies, it may be more difficult to establish.