Mixed Strategy Equilibrium and Deep Covering in Multidimensional Electoral Competition

John Duggan[*] Matthew O. Jackson[†]

Preliminary: January 13, 2004
This Draft: December 12, 2005

Abstract

We prove existence of mixed strategy equilibrium in a general model of elections that includes multicandidate Downsian models and probabilistic voting models as special cases. We do so by modelling voters explicitly as players, enabling us to resolve discontinuities in the game between the candidates, which have proved a barrier to existence. We then give a partial characterization: the supports of equilibrium mixed strategies in two-candidate Downsian games must lie in the deep uncovered set.

Journal of Economic Literature Classification Numbers: C72, D72.

Key words: Electoral Competition, Colonel Blotto, elections, mixed strategy equilibrium, undominated strategies, uncovered set

[*]Department of Political Science and Department of Economics, University of Rochester, Rochester, NY 14627, http://www.rochester.edu/College/PSC/duggan, email: dugg@troi.cc.rochester.edu. Financial support under NSF grant SES-0213738 is gratefully acknowledged.

[†]Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, http://www.hss.caltech.edu/~jacksonm/Jackson.html, email: jacksonm@hss.caltech.edu. Financial support from the Center for Advanced Studies in the Behavioral Sciences, the Guggenheim Foundation and from the NSF under grant SES-0316493 is gratefully acknowledged.

1 Introduction

The formal theory of elections, originated by Hotelling (1929), Downs (1957), and Black (1958), has developed despite a well-known and unresolved foundational issue arising in the basic model: the existence of equilibrium in multidimensional policy spaces. Plott (1967) showed that a multidimensional model typically does not admit a pure strategy equilibrium, and in most cases it has remained unclear even whether mixed strategy equilibria exist. As a consequence, much research has either been limited to unidimensional policy spaces, or has been confined to special variations of the basic model that circumvent the existence problems; such as probabilistic voting, policy-motivated candidates, citizen-candidates, candidates who wish to maximize vote share from a heterogeneous continuum of voters, and repeated elections.

Clearly multidimensional policy spaces are fundamental to the understanding of political interaction, and it is essential to be able to model them. Moreover, it would be especially useful to incorporate multidimensional policies in classical models. The hurdle that such models face, and which makes the existence questions so hard to answer, is that such models exhibit discontinuities that could potentially preclude the existence of equilibrium (even when mixed strategies are allowed). To see this, consider the most basic Downsian model of competition between two candidates in a multidimensional space. Each candidate simultaneously chooses a policy from some multi-dimensional space. Voters then vote for their most preferred candidate. A candidate receives a payoff of 1 if her policy position is preferred by a majority of voters to the opposing candidate's position and a payoff of -1 if this situation is reversed (and a payoff of 0 otherwise). There are discontinuities in candidates' payoffs, since when they choose the same policies (or, more generally, any time they choose policies that render some voters indifferent), the voter behavior changes from preferring one of the candidates outright and voting for that candidate, to being indifferent. Thus, a candidate can have a majority of voters all along a sequence of policy choices, and not have a majority in the limit; or vice versa. This means that standard results on existence of mixed strategy equilibria do not apply. This problem is not only endemic to Downsian models, but to a very wide variety of other models as well.

In this paper, we establish existence of equilibria in a general class of electoral games, including the Downsian model and others where such discontinuities are present. Our starting point is the observation that the discontinuities of the classical Downsian game, as well as many other voting games, arise from the implicitly fixed behavior of indifferent voters. It is generally assumed that a voter who is indifferent between two candidates flips a fair coin. The central observation that allows us to overcome these difficulties is that by including the voters as players in the game, so that an indifferent voter might flip a biased coin or simply vote for one candidate or the other when indifferent, existence can be established in a very wide class of games, including all of the classical models.

Our proof technique builds on an insightful theorem by Simon and Zame (1990) which shows that existence of equilibrium in games with discontinuities can be overcome if outcomes at discontinuity points can be properly selected. We show that by including voters as players in the game, the necessary selections of election outcomes can in fact be rationalized by equilibrium play on the part of voters.[1]

Our results apply to a general electoral model, which includes complete information as a special case, but also allows voters to have uncertain preferences. Candidates can care about policies, or winning, or vote-shares, or elaborate combinations of these outcomes. Similarly, voters can care about who wins, or which policy is chosen, or some combination of the two. The candidates strategies are general sets of campaign strategies, which we interpret as campaign platforms. The strategy sets are sufficiently general that a strategy could be a plan for running a campaign in the reduced form of a highly complex election: we are not bound by the traditional assumption of simultaneous chosen policy platforms. We eliminate dominated voting strategies, but in contrast to the classical approach, we model the votes of indifferent voters as endogenous. We prove the existence of equilibria in

[1] It remains an open question to what extent equilibria exist when voter behavior is fixed (to say flipping fair coins when indifferent). Theorems on equilibrium existence in discontinuous games have proved difficult to apply because they employ complex sufficient conditions, e.g., Dasgupta and Maskin's (1986) weak upper semicontinuity or, more generally, Reny's (1999) better reply security. Duggan (2005) provides a sufficient condition for better reply security and verifies that it holds in the classical Downsian model with three voters, but his approach does not extend to multiple voters. Another approach from the auctions literature (e.g., Lebrun (200?) and Jackson and Swinkels (2005)) that might be fruitful is to prove existence with endogenous outcomes, and then show that ties never occur and that changing voter behavior at ties would not matter.

the multicandidate model – indeed, we establish existence of Duvergerian equilibria in which only two candidates receive votes.

Beyond the results on existence, we also show that candidates only choose strategies that lie in a refinement of the top cycle of the majority voting relation. In particular, we show that the support of winning candidates' strategies lie in a variation of the uncovered set. This shows that results found to hold in finite policy settings have proper analogs in continuous multidimensional policy settings. In particular, Laffond, Laslier, and Le Breton (1993) showed that if policies must be chosen from a finite set and there are no majority indifferences, then the mixed strategy equilibrium of a Downsian game is unique and that the support of the equilibrium strategy lies in the uncovered set.[2] McKelvey (1986), in the context of the multidimensional spatial model, conjectured that the support of equilibrium mixed strategies - *if any exist* – must lie in the uncovered set, and this was confirmed by Banks, Duggan, and Le Breton (2002). These results leave open the issue of existence in the classical model, and also whether the equilibrium strategies must have this property outside of a particular spatial model. Here we prove existence and then also show that the support of equilibrium strategies is in the uncovered set in a general class of elections.

The paper proceeds as follows. Section 2 presents an preview of the general electoral model and the results of the paper. Section 3 provides a full description of the setting and the details omitted in the overview. Section 4 contains equilibrium existence results that "black box" the behavior of voters as a correspondence satisfying several technical properties. In Section 5, we apply these results to establish existence in the multicandidate complete information model. In Section 6, we develop the deep uncovered set and show that this set bounds the equilibrium outcomes of two-candidate Downsian elections. Section 7 applies our results on endogenous voting equilibria to a general probabilistic voting model. Our final section discusses an extension to a world where public randomization is possible. The proofs appear in an appendix, and another appendix contains an example illustrating the tightness of our bounds on equilibrium strategies.

[2] In the presence of majority indifferences, the literature assumed that indifferent voters randomize over candidates with equal probabilities. Assuming a finite policy space, Dutta and Laslier (1999) show that supports of equilibrium strategies still lie in the uncovered set.

2 A Preview of the Model and Results

In this section, we provide statements of special cases of our electoral model and results. This section lacks the full generality of the model and results that follow in later sections, but allows for a simpler presentation that covers many of the standard cases from the literature and should be more easily accessed by the non-specialist.

We consider an election among a set of candidates $\mathcal{C} = \{A, B, \ldots, M\}$ to be decided by an electorate of voters $\mathcal{V} = \{1, 2, \ldots, n\}$ using plurality rule.

The candidates $C \in \mathcal{C}$ simultaneously choose policy positions, x_C, in a compact subset X of $\mathbb{R}^\ell$. Let $\overline{x} = (x_A, \ldots, x_M)$ denote a profile of policy platforms for the candidates, and let $\overline{X}$ denote the set of platform profiles.

Voters observe the policy positions and then cast votes. Voter i, having observed the candidates' platforms, casts a ballot b_i in $\mathcal{C}$ for one of the candidates, with no abstention allowed. Let $\overline{b} = (b_1, \ldots, b_n) \in \mathcal{C}^n$ denote a profile of ballots, i.e., an election return which indicates which candidate each voter voted for. The winner is the candidate with the greatest number of votes, with ties broken by an even lottery among the tied candidates. Let

$$W(\overline{b}) = \left\{ C \in \mathcal{C} \; : \; \#\{i : b_i = C\} \geq \#\{i : b_i = C'\} \text{ for all } C' \in \mathcal{C} \right\}$$

denote the set of potential winners (before any ties are broken), for a given election return ($\overline{b}$.

Candidate C has a continuous utility function $u_C : \mathcal{C}^n \times \overline{X} \to \mathbb{R}$, where the payoff to candidate C from an outcome $(\overline{b}, \overline{x})$ is $u_C(\overline{b}, \overline{x})$. For our existence results, we impose a restriction on the ways in which u_C is allowed to depend on election returns. Because that restriction is mainly technical, here we will simply note that it is satisfied in the following three special cases, all prominent in the literature, and refer the interested reader to later sections for the fuller model.

- win-motivation:

$$u_C(\overline{b}, \overline{x}) \;=\; \begin{cases} \frac{1}{\#W(\overline{b})} & \text{if } C \in W(\overline{b}) \\ 0 & \text{else.} \end{cases}$$

- vote-motivation:

$$u_C(\bar{b}, \overline{x}) \;\; = \;\; \#\{i \; : \; b_i = C\}.$$

- policy-motivation:

$$u_C(\bar{b}, \overline{x}) \;\; = \;\; \sum_{C' \in W(\bar{b})} \frac{u_C(x_{C'})}{\#W(\bar{b})}.$$

Voter i with a continuous utility function $u_i \colon \mathcal{C} \times X \to I\!R$, and the payoff to voter i from outcome $(\bar{b}, \overline{x})$ is

$$u_i(\bar{b}, \overline{x}) \;\; = \;\; \sum_{C' \in W(\bar{b})} \frac{u_i(C', x_{C'})}{\#W(\bar{b})}.$$

An important special case is where voters care only about policy outcomes: we say that voter i is *policy-oriented* if $u_i(C, x)$ does not depend on C.

We allow candidates and voters to use mixed strategies in equilibrium, as in many applications this is actually necessary for equilibrium existence. Some of our formal results focus on electoral equilibria in which candidate mixed strategies are *symmetric,* i.e., the candidates mix over policy platforms with the same probability distribution. We consider voting equilibria in which no voter puts positive probability on a "dominated" strategy. For technical reasons (discussed in the next section), our refinement is weaker than the usual notion of undominated strategies: we say a ballot b_i for voter i is *undominated** if b_i is not the strictly worst candidate for voter i. Our definition of undominated strategies coincides with the usual one when there are only two candidates.

We say that an electoral equilibrium is *Duvergerian* if there are two candidates such that, for all profiles of platforms, only those two candidates receive votes. That is, there are only two "viable" candidates. In this case, we say a voting strategy for i is *symmetric* if i's vote between the two viable candidates depends not on the name of the candidate, but only on the candidates' platforms. So, for example, if i votes for candidate A with probability $1/3$ and candidate B with probability $2/3$, and then we consider a platform profile in which the candidates' platforms are interchanged, then the voter must now vote for A with probability $2/3$ and for B with probability $1/3$.

Our first result establishes the general existence of subgame perfect equilibria in the electoral model. (Theorems are numbered as they are presented in later sections.) The key step in the proof of Theorem 3 is to convert the electoral game to a game with an "endogenous sharing rule," allowing us to apply results of Simon and Zame (1990) and Jackson, Simon, Swinkels, and Zame (2002) to specify the behavior of indifferent voters appropriately.

THEOREM **3** *Consider the multi-candidate complete information electoral model, where there are at least four voters or there are only two candidates. There exists a subgame perfect equilibrium in undominated* voting strategies. Moreover, there exists such an equilibrium that is Duvergerian, in which two arbitrary candidates A and B receive votes. If all voters are policy-oriented and candidates A and B are both win-motivated or both vote-motivated, then there exists such an equilibrium that is symmetric on $\{A, B\}$.*

We next give a partial characterization of equilibrium policy platforms in the Downsian model, where we assume candidates seek to maximize their probabilities of winning. We show that policy platforms lie in the deep uncovered set, a centrally located region of the policy space that has received considerable attention in the political science literature. It is one of many possible extensions of the uncovered set to settings where ties are possible in the majority voting relationship.

The deep uncovered set consists of the policies x such that, for every other policy y, either x is weakly majority-preferred to y or there is a policy z such that x is weakly majority-preferred to z and z is weakly majority-preferred to y.

THEOREM **5** *In the two-candidate Downsian model (where candidates are win motivated and voters are policy oriented), in every subgame perfect equilibrium that has symmetric and undominated* voting strategies, the support of each candidate's strategy lies in the deep uncovered set.*

Theorem 4 tells us that every policy in the support of an equilibrium mixed strategy is weakly majority-preferred to every other policy in at most two steps.

We also provide an existence result for a general model of probabilistic voting in two-candidate elections. Each voter's policy preferences now depend on a type t_i, unobserved by the candidates, and we extend voter utilities to include this as an argument, as in $u_i(C, x, t_i)$. We do not restrict voter preferences, thereby capturing two special cases prominent in the electoral modelling literature.

- stochastic candidate bias:

$$u_i(C, x, t_i) = u_i(x) + t_{i,C}.$$

- stochastic policy preference:

$$u_i(C, x, t_i) = u_i(x, t_i).$$

Although most applications of the probabilistic voting model impose some structure on the distribution of voter types (requiring, for instance, that voters are indifferent with probability 0), we do not restrict the distribution of voter types in any way.

THEOREM **6** *Consider the multi-candidate probabilistic voting model, where there are at least four voters or there are only two candidates. There exists a perfect Bayesian equilibrium in undominated* voting strategies. Moreover, there exists such an equilibrium that is Duvergerian, in which two arbitrary candidates A and B receive votes. If all voters are policy-oriented and candidates A and B are both win-motivated or both vote-motivated, then there exists such an equilibrium that is symmetric on $\{A, B\}$.*

Note that the complete information model is a special case of the probabilistic voting model.

Thus, we prove existence of mixed strategy equilibria the canonical models of elections, and we give a partial characterization in terms of the uncovered set for the basic two-candidate Downsian model. In the sequel, we also give convergence results as we vary voter preferences and (in the probabilistic voting model) the distribution of voter types.

3 A General Model of Elections

We now present the full model.

3.1 Candidates and Policies

Each candidate $C \in \mathcal{C}$ chooses a policy position, x_C, in a second countable, locally compact Hausdorff space X_C. Let $\overline{x} = (x_A, \ldots, x_M)$ denote a profile of policy platforms for the candidates, and let $\overline{X} = \prod_{C \in \mathsf{C}} X_C$ denote the set of platform profiles. Let $X = \bigcup X_C$ denote the space of all possible policies. A mixed platform for a candidate $C \in \mathcal{C}$ is a Borel probability measure ξ_C on X_C, with the product measure $\xi = \xi_A \times \cdots \times \xi_M$ representing a profile of mixed platforms for the candidates.

3.2 Voting Correspondences

We explicitly model voters' behavior in subsequent sections, but here we analyze the game among just among the candidates, black-boxing voter behavior. For now, we capture voter behavior through a correspondence, $\Psi \colon X^M \rightrightarrows \Delta(\mathcal{C}^n)$, from platform profiles to probability distributions over election returns. The correspondence Ψ represents possible (equilibrium) voter behavior contingent on the candidates' positions. A measurable selection ψ from Ψ determines a well-defined (but generally discontinuous) game among the candidates.

3.3 Election Returns and Statistics

We wish to admit applications where candidates have preferences over more than simply whether they win and/or which policy is enacted. For instance, we wish to allow for applications where candidates care about their vote totals or vote shares. In order to define candidate payoffs in such a general manner, we let candidate preferences over election returns be dependent on a statistic $\sigma \colon \mathcal{C}^n \to S$, where S is a nonempty, compact, and convex subset of a finite-dimensional Euclidean space.

We impose two conditions on this statistic, both of which only apply to situations where at most two candidates receive votes. Let $\overline{B}^{AB}$ denote the set of election returns in which only candidates A and B receive votes; i.e.,

$$\overline{B}^{AB} \quad = \quad \left\{ \bar{b} \in \mathcal{C}^n \mid \text{ for all } i \in \mathcal{V},\, b_i \in \{A, B\}. \right\}$$

The first condition that we impose on the statistic is that whenever ballots are restricted to any pair of candidates, the range of statistic is one-dimensional ($\{\sigma(\bar{b}) | \bar{b} \in \overline{B}^{AB}\}$ is one-dimensional for any A and B).

The second condition that we impose on the statistic is that whenever ballots are restricted to any pair of candidates, the statistic is nondecreasing in the set of votes cast for one of the candidates (and hence nonincreasing in the votes cast for the other). Given two candidates A and B and two election returns $\bar{b}, \bar{b}' \in \mathcal{C}^n$, write $\bar{b} \geq_A \bar{b}'$ if everyone who voted for A in $\bar{b}$ also votes for A in $\bar{b}'$, i.e., for all $i \in \mathcal{V}$, $b'_i = A$ implies $b_i = A$. Then the monotonicity condition is for all $A, B \in \mathcal{C}$, either

$$\text{for all } \bar{b}, \bar{b}' \in \overline{B}^{AB},\ \bar{b} \geq_A \bar{b}' \text{ implies } \sigma(\bar{b}) \geq \sigma(\bar{b}'),$$

or

$$\text{for all } \bar{b}, \bar{b}' \in \overline{B}^{AB},\ \bar{b} \geq_A \bar{b}' \text{ implies } \sigma(\bar{b}) \leq \sigma(\bar{b}').$$

3.4 Candidate Preferences

Each candidate C has a utility function $u_C : S \times \overline{X} \to I\!R$, which is jointly continuous. Moreover, $u_C(s, x)$ is affine in s. The payoff to candidate C from an outcome $(\bar{b}, \bar{x})$ is $u_C(\sigma(\bar{b}), \bar{x})$. Let u denote a profile of utility functions, one for each candidate. An *electoral game* is then a triple $(\overline{X}, \Psi, u)$ satisfying the maintained assumptions stated above.

The main approaches to modeling candidates in the literature all satisfy our assumptions. Some prominent special cases of the model are:

- win-motivation: let S be the unit simplex in $I\!R^M$, let

$$\sigma_C(\bar{b}) \quad = \quad \begin{cases} \frac{1}{\#W(\bar{b})} & \text{if } C \in W(\bar{b}) \\ 0 & \text{else,} \end{cases}$$

and define $u_C(s, \overline{x}) = s_C$. That is, candidate C's payoff is just the candidate's probability of winning.

- policy-motivation: let S and σ be as above, but let

$$u_C(s, \overline{x}) = \sum_{C' \in \mathsf{C}} s_C \hat{u}_C(x_{C'}),$$

where $\hat{u}_C \colon X \to \mathbb{R}$ is continuous. That is, the candidate's payoff is the expected utility generated by a utility function defined over policy outcomes.

- vote-motivation: let S be the cube $[0, n]^M$ in $\mathbb{R}^M$, let

$$\sigma_C(\overline{b}) = \#\{i : b_i = C\} - \max_{C' \neq C} \#\{i : b_i = C'\},$$

and define $u_C(s, \overline{x}) = s_C$. That is, the candidate's payoff is just the number of votes received in the election.[3]

To verify the dimensionality assumption for win-motivated and policy-motivated candidates, just note that if two candidates A and B receive all of the votes in an election return $\overline{b}$, then $\sigma_A(\overline{b}) + \sigma_B(\overline{b}) = 1$; and under vote-motivation, $\sigma_A(\overline{b}) + \sigma_B(\overline{b}) = n$. Our monotonicity assumption is clearly satisfied by all of the above specifications.

Note that the distinction between win-motivation and vote-motivation is often overlooked, as these two objectives lead to the same pure strategy equilibria. When we consider equilibria in mixed strategies, however, the distinction becomes potentially significant. The same observation holds when voting behavior is probabilistic, and the distinction between the two objectives is, accordingly, well-known in that literature.

We extend payoffs to lotteries over outcomes in the standard expected utility manner, thereby inducing preferences over strategy profiles for candidates. Given a measurable selection ψ of voting behavior, preferences of a

[3] It could also be that the candidate cares about vote share, or vote differences. For instance, the condition is met if $\sigma_C(\overline{b}) = \#\{i : b_i = C\} - \max_{C' \neq C} \#\{i : b_i = C'\}$, or $\sigma_C(\overline{b}) = \#\{i : b_i = C\} / \max_{C' \neq C} \#\{i : b_i = C'\}$, or $\sigma_C(\overline{b}) = \#\{i : b_i = C\}/n$, etc.

candidate C over platform profiles are represented by the measurable payoff function $u_C^\psi \colon \overline{X} \to I\!R$, defined by

$$u_C^\psi(\overline{x}) = \sum_{\overline{b} \in \mathsf{C}^n} u_C(\sigma(\overline{b}), \overline{x})\, \psi(\overline{x})(\overline{b}).$$

Given the selection ψ, the expected payoff of candidate C from the profile ξ of mixed platforms is then

$$U_C^\psi(\xi) = \int u_C^\psi(\overline{x})\, \xi(d\overline{x}).$$

3.5 Equilibrium

An *endogenous voting equilibrium* is a measurable selection ψ from Ψ, and a profile of mixed strategies ξ, such that ξ forms a Nash equilibrium under the payoff functions U_C^ψ.

To define symmetric equilibrium, we use the following notational conventions. Given two candidates A and B and platform profile $\overline{x}$, let $\overline{x}^{AB}$ denote the result of interchanging the two candidates' platforms, leaving the other candidates unchanged, i.e.,

$$x_C^{AB} = \begin{cases} x_B & \text{if } C = A \\ x_A & \text{if } C = B \\ x_C & \text{else.} \end{cases}$$

Similarly, let $\overline{b}^{AB}$ denote the result of interchanging ballots for the two candidates, i.e., for all $i \in \mathcal{V}$,

$$b_i^{AB} = \begin{cases} B & \text{if } b_i = A \\ A & \text{if } b_i = B \\ b_i & \text{else.} \end{cases}$$

If $X_C = X$ for all candidates C, then we say the selection ψ is *symmetric* on $\mathcal{C}' \subseteq \mathcal{C}$ if

- for all $\overline{x} \in \overline{X}$, $\psi(\overline{x})$ has support in $\mathcal{C}'$, i.e., $\psi(\overline{x})(\overline{b}) > 0$ implies that for all $i \in \mathcal{V}$, $b_i \in \mathcal{C}'$,

- ψ is symmetric in the positions of candidates in $\mathcal{C}'$, i.e., for all $A, B \in \mathcal{C}'$, all $\overline{x} \in \overline{X}$, and all $\bar{b} \in \mathcal{C}^n$, we have $\psi(\overline{x})(\bar{b}) = \psi(\overline{x}^{AB})(\bar{b}^{AB})$.

We say ξ is *symmetric on* $\mathcal{C}'$ if for all $A, B \in \mathcal{C}'$, $\xi_A = \xi_B$. Finally, we say an endogenous voting equilibrium (ξ, ψ) is *symmetric on* $\mathcal{C}'$ if ξ and ψ are both symmetric on $\mathcal{C}'$. That is, all candidates in $\mathcal{C}'$ use the same mixed strategy, only those candidates receive votes, and voters' ballots depend on the platforms of those candidates in a neutral way. When $\mathcal{C}' = \mathcal{C}$, we use the above terminology without reference to the subset $\mathcal{C}'$.

The electoral game $(\overline{X}, \Psi, u)$ is *symmetric on* $\mathcal{C}' \subseteq \mathcal{C}$ if there is a measurable selection ψ that is symmetric on $\mathcal{C}'$ and for all $A, B \in \mathcal{C}'$ and all $\overline{x} \in \overline{X}$, we have $u_A^\psi(\overline{x}) = u_B^\psi(\overline{x}^{AB})$. That is, the game is symmetric on a subset of candidates if there is a selection that is symmetric on that subset and, given that selection, the payoffs of those candidates are symmetric.

4 Endogenous Voting Equilibria

We now present results on existence and continuity of endogenous voting equilibria that only rely on the structure provided above. In the following sections, we then use these results to derive further results for electoral models where voters are players.

Our first result establishes the existence of an endogenous voting equilibrium.

Given a distribution $\nu \in \Delta(\mathcal{C}^n)$, let $\sigma\Psi \colon \overline{X} \rightrightarrows S$ denote the correspondence defined by

$$\left\{ \int \sigma(\bar{b})\nu(d\bar{b}) \mid \nu \in \Psi(\overline{x}) \right\}.$$

That is, $\sigma\Psi(\overline{x})$ consists of the mean of the statistic generated by all possible voting behavior at platform vector $\overline{x}$.

THEOREM **1** *Suppose that X is a compact metric space and that $\sigma\Psi$ is upper hemi-continuous with nonempty, compact, and convex values. Then there* *exists a Borel measurable selection ψ from Ψ and a profile ξ of mixed strategies, such that (ξ, ψ) forms an endogenous voting equilibrium. Moreover, if the electoral game is symmetric on some subset of candidates $\mathcal{C}'$, then there exists an endogenous voting equilibrium (ξ, ψ) that is symmetric on $\mathcal{C}'$.*

The theorem follows from a direct application of theorems by Simon and Zame (1990) to establish existence, and by Jackson, Simon, Swinkels and Zame (2002) to obtain existence of symmetric equilibria. The proof (which appears in Appendix A) consists of verifying that the conditions required by those theorems are satisfied in our setting.

We remark that one can choose the strategies of the candidates who are outside of the support of ψ arbitrarily, and so the candidates' strategies can be selected to be fully symmetric, even when ψ has support only on some subset of candidates.

Another useful result is a convergence result across electoral games. This is useful for working with finite approximations, as well as deriving results about existence of perfect equilibrium, among other things. Fix the set of candidates. We say that a sequence of electoral games $(\overline{X}^r, \Psi^r, u^r)$ converges to $(\overline{X}^0, \Psi^0, u^0)$ if:

- for each $C \in \mathcal{C}$, each X_C^r and X_C^0 are compact metric spaces that are subsets of some common compact metric space X_C, and such that X_C^r converges to X_C^0 in the Hausdorff metric,
- the graph of Ψ^r converges to the graph of Ψ^0 in the Hausdorff metric,
- for every $\varepsilon > 0$ there exists $\delta > 0$ and r such that, for all $r' > r$, all $C, C' \in \mathcal{C}$, and all $\overline{x}^r \in \overline{X}^r$ and $\overline{x}^0 \in \overline{X}^0$ such that $d(\overline{x}^r, \overline{x}^0) < \delta$, we have $|u_C^r(C', \overline{x}^r) - u_C^0(C', \overline{x}^0)| < \varepsilon$.

Our next result shows that limits of equilibria corresponding to such a sequence of electoral games are indeed equilibria of the limiting game. The result follows directly from Theorem 2 in Jackson, Simon, Swinkels, and Zame (2002) (noting from the proof of Theorem 1 that the necessary conditions to apply their Theorem 2 are satisfied).

THEOREM **2** *Consider a sequence of electoral games* $(\overline{X}^r, \Psi^r, u^r)$ *that converge to* $(\overline{X}^0, \Psi^0, u^0)$*, and a corresponding sequence of endogenous voting equilibria* (ξ^r, ψ^r)*. Then there exist an endogenous voting equilibrium* (ξ^0, ψ^0) *of* (X^0, Ψ^0, u^0) *and a subsequence* $(\xi^{r'}, \psi^{r'})$ *such that*

- $\xi_C^{r'} \to \xi_C^0$ *weak*,*
- $\psi^{r'}\xi^{r'} \to \psi^0\xi^0$ *weak*,*[4] *and*
- *for each* $C \in \mathcal{C}$*,* $U_C^{r',\psi^{r'}}(\xi^{r'}) \to U_C^{0,\psi^0}(\xi^0)$*.*

5 Equilibrium with Voters as Players

We now explicitly model voters as part of the game, and show that we can still establish existence of equilibrium. As mentioned in the introduction, the key to establishing this is to account for how voters vote when they are indifferent, and noting that this may have to differ from the flip of a fair coin in order to guarantee existence.

5.1 Voters' Preferences and Strategies

Now let X be a compact metric space, and let each voter i have a continuous utility function $u_i\colon \mathcal{C} \times X \to I\!R$. The payoff to voter i from outcome $(\bar{b}, \overline{x})$ is

$$u_i(\bar{b}, \overline{x}) \quad = \sum_{C' \in W(\bar{b})} \frac{u_i(C', x_{C'})}{\#W(\bar{b})}.$$

Thus, voters care about who wins the election and/or what the policy outcome is. They voters do not care directly for the specific ballot they cast except to the extent that it determines the outcome of the election.

[4] $\psi\xi$ denotes the set function defined by

$$\psi\xi(Y) \quad = \quad \int_Y \psi(\overline{x})\xi(d\overline{x}),$$

which measures the ballots cast when candidates use platform profiles in Y.

An important special case of the model is that in which voters are *policy-oriented,* i.e., for each i, $u_i(C,x)$ is independent of the first argument and depends only on the policy.

A strategy for voter i is a behavioral strategy β_i, where $\beta_i(\cdot|\overline{x})$ is a probability distribution on $\mathcal{C}$, expressing the probability that voter i votes for each candidate given that the candidates have chosen policy positions $\overline{x}$. Viewing $\beta_i(\cdot|\overline{x})$ as a vector in the simplex, it is required to be Borel measurable as a function of $\overline{x}$. Let the product measure $\beta(\cdot|\overline{x}) = \beta_1(\cdot|\overline{x}) \times \cdots \times \beta_n(\cdot|\overline{x})$ represent a list of behavioral strategies for the voters, and let $u_C^{\beta}(\overline{x})$ denote the expected payoff to candidate C when the candidates take positions $\overline{x}$ and the voters employ strategies β. That is,

$$u_C^{\beta}(\overline{x}) \;\; = \;\; \sum_{\bar{b}\in \mathbf{C}^n} u_C(\sigma(\bar{b}),\overline{x})\beta(\bar{b}|\overline{x}).$$

Let $U_C^{\beta}(\xi)$ denote the expected payoff to candidate C when voting behavior is given by β and candidates use

the mixed platforms ξ. That is,

$$U_C^{\beta}(\overline{x}) \;\; = \;\; \int u_C^{\beta}(\overline{x})\xi(d\overline{x}).$$

We seek a profile (ξ,β) of candidate and voter strategies that form a subgame equilibrium of the electoral game, after suitable refinements are imposed.

5.2 Refining to Undominated* Strategies

A known drawback of Nash-based equilibrium concepts in voting games is that a wide variety of (degenerate) outcomes can be supported by a Nash equilibrium of the voting game by specifying voting strategies that make no voters pivotal. A standard approach to dealing with this problem is to require that voters not use dominated strategies. Say that β_i is *undominated* if for all platform profiles $\overline{x}$, $\beta_i(C|\overline{x}) > 0$ implies that either there exists $C' \neq C$ such that $u_i(C',x_{C'}) < u_i(C,x_C)$ or for all C' we have $u_i(C,x_C) = u_i(C',x_{C'})$.

In discontinuous games with continuum action spaces, however, it is generally natural (in order to guarantee existence) only to require that voters use strategies that are in the closure of the set of undominated strategies.[5] We use a slightly weaker definition, saying that a behavioral strategy β_i is *undominated** if for all platforms $\overline{x}$, $\beta_i(C|\overline{x}) > 0$ implies that either there exists $C' \neq C$ such that $u_i(C, x_C) \geq u_i(C', x_{C'})$.[6] Note that when there are just two candidates, the undominated* voting strategies coincide with the voting strategies that are undominated in the conventional sense.

When voters are policy-oriented, it is natural to look for equilibria where voters treat the set of viable candidates symmetrically. Given a subset $\mathcal{C}'$ of candidates, we say that β is *symmetric on* $\mathcal{C}'$ if for all $A, B \in \mathcal{C}'$ and all $\overline{x} \in \overline{X}$,

$$\beta_i(A|\,\overline{x}) \;\;=\;\; \beta_i(B|\,\overline{x}^{AB}).$$

In particular, if there are just two viable candidates, say A and B, and these candidates take the same policy position x, then all voters emplying symmetric strategies would flip fair coins to decide their votes: $\beta_i(A|x, x) = \beta_i(B|x, x) = \frac{1}{2}$.

We now consider existence in the general multi-candidate model, allowing for an arbitrary number of candidates, for arbitrary candidate motivations, and for voters who care not just about policy but also for the particular candidates elected. We establish existence of a subgame perfect equilibrium in undominated* voting strategies. Moreover, we prove existence of an equilibrium in which there are just two candidates who receives votes in all subgames. Formally, we say β is *Duvergerian* if there are two candidates $A, B \in \mathcal{C}$ such that for all $\overline{x} \in \overline{X}$ and all voters i, $\beta_i(\{A, B\}|\overline{x}) = 1$. When voters are policy-oriented, we can find an equilibrium where voting strategies

[5] For instance, see Jackson and Swinkels (2005) and the discussion there.

[6] Undominated strategies are those where a voter only votes for a least preferred candidate in the case where the voter is totally indifferent. Here, undominated* strategies have a voter only voting for a least preferred candidate in the case where the voter is indifferent between that candidate and at least one other. The idea here is that a vote v_i is undominated* at $\overline{x}$ if there is some sequence $\overline{x}^k \to x$ such that the vote v_i is undominated for each $\overline{x}^k$. Provided that preferences are such that whenever a voter finds some candidate pair tied at a given policy profile there exists arbitrarily small perturbations in policies that can lead the voter to prefer either candidate, and that either there are at least four voters or there are just two candidates, then this notion of closure leads to our formal definition of undominated* strategies. In order not to add such additional preference restrictions, we simply define undominated* directly, rather than through the closure.

are symmetric on those two candidates.[7]

THEOREM **3** *Consider the multi-candidate complete information electoral model, where there are at least four voters or there are only two candidates. There exists a subgame perfect equilibrium in undominated* voting strategies. Moreover, there exists such an equilibrium that is Duvergerian, in which two arbitrary candidates A and B receive votes. If all voters are policy-oriented and candidates A and B are both win-motivated or both vote-motivated, then there exists such an equilibrium that is symmetric on $\{A, B\}$.*

6 Downsian Elections and the Deep Uncovered Set

In this section, we consider a special case of the multicandidate complete information model. That special case, the *Downsian model,* is one where candidates are win-motivated and voters are policy-oriented.

Theorem 3 directly implies the existence of a subgame perfect equilibrium in which only two candidates, say A and B, receive votes, candidate strategies are symmetric, and voting strategies are symmetric on $\{A, B\}$ and undominated*. In such equilibria, the positions of candidates other than A and B are immaterial and so, without loss of generality, we focus attention on two-candidate elections. That is, in this section let $\mathcal{C} = \{A, B\}$. Here undominated and undominated* voting strategies coincide.

COROLLARY **1** *In the two-candidate Downsian model, there exists a symmetric subgame perfect equilibrium in undominated voting strategies.*

6.1 The Deep Uncovered Set

While existence of equilibrium follows from Theorem 3, we have not yet addressed the characterization of equilibrium platforms. We do so now by

[7] Note that this is not the same as saying that voters always flip a coin when they are indifferent. It only requires that they break ties in symmetric ways more generally (e.g., if they are indifferent between x and y and vote for A when A proposes x and B proposes Y, then they do the reverse when the candidates' platforms are reversed).

demonstrating that in the equilibrium of Corollary 1, the candidates adopt platforms in the "uncovered set" (a prominent choice set in the extensive literature on voting rules) with probability one.

In settings with a continuum of outcomes, there are various definitions of the uncovered set that differ with respect to treatments of ties under the majority relation. So, we need to be explicit in defining the set that we consider.

Define strict and weak majority preferences, respectively, as follows: for all $x, y \in X$,

$$xPy \;\Leftrightarrow\; |\{i \mid u_i(x) > u_i(y)\}| > \frac{n}{2}$$
$$xRy \;\Leftrightarrow\; |\{i \mid u_i(x) \geq u_i(y)\}| \geq \frac{n}{2}.$$

The relations P and R are dual, in the sense that xPy if and only if not yRx (and xRy if and only if not yPx). Clearly, P is asymmetric and R is complete, and our continuity assumption on u_i implies that P is open and R is closed. We use the standard notation,

$$P(x) = \{y \in X \mid yPx\} \quad \text{and} \quad R(x) = \{y \in X \mid yRx\},$$

for the upper sections of these relations. As is customary, let xIy denote majority indifference, where neither xPy nor yPx (or equivalently, xRy and yRx).

We bound equilibrium policy positions by a weak version of the uncovered set, derived as the maximal elements of a strong version of the usual covering relation. Following Duggan (2006), we say that x *deeply covers* y, written $x\,DC\,y$, if $R(x) \subseteq P(y)$. That is, x deeply covers y if every alternative weakly majority-preferred to x is strictly majority-preferred to y. Equivalently, x deeply covers y if whenever y weakly defeats some alternative z, then x strictly defeats z. Note that this implies that if $x\,DC\,y$ implies xPy, i.e., deep covering is a subrelation of P. Obviously, DC is transitive. Moreover, Duggan (2006) shows that the deep covering relation, DC, is open under our maintained assumptions. The *deep uncovered set* is defined as the maximal elements of the deep covering relation:

$$DC \;=\; \{x \in X \mid \text{ there is no } y \in X \text{ such that } y\,DC\,x\}.$$

Assuming X is compact, the fact that DC is transitive and open immediately yields non-emptiness of the deep uncovered set, and in fact DC is compact. Note that the deep uncovered set is equivalently defined by the following "two-step" principle:

$$x \in DC \quad \Leftrightarrow \quad \text{for all } y \in X, \text{ there exists } z \in X \text{ such that } xRzRy.$$

Thus, a policy position x is in the deep uncovered set if and only if, for every other position y, x is either weakly majority-preferred to y directly or indirectly, via some other alternative.

The next theorem shows that, given any equilibrium in which voters use symmetric, undominated strategies, the candidates locate in the deep uncovered set with probability one. Note that the result does not assume candidate strategies are symmetric, so the result holds even for non-symmetric equilibria.

Theorem 4 *In the two-candidate Downsian model, let (ξ, β) be a subgame perfect equilibrium in symmetric, undominated voting strategies. Then the supports of both candidates' strategies lie in the deep uncovered set; that is, $\xi_A(DC) = \xi_B(DC) = 1$.*

Although Theorem 4 applies to equilibria in which candidates may use non-symmetric strategies, the focus on equilibria where voters are using symmetric strategies is needed for the result. For instance, if voters always favored one candidate in the case of ties, it is easy to construct examples where that candidate wins and the other candidate chooses arbitrary policies. In Appendix B, we present a (finite) example showing that if voters favor one candidate over the other, then there exist equilibria in which one candidate locates outside the weak uncovered and that candidate wins with positive probability.

7 Probabilistic Voting

We now address another prominent model from the literature, namely where voters' preferences are unknown to the candidates. The candidates have beliefs about voter preferences, so that given any pair of campaign platforms, any voter's behavior is a random variable from the perspective of the candidates. Here we examine a general setting that captures models of "probabilistic voting" from the political science literature as special cases. We return to the general case in terms of candidate and voter preferences.

Define the *probabilistic voting model* as follows. The policy space, X, is a compact metric space. For each voter i let T_i be a space of types, with a generic type for voter i denoted t_i. Let $T = \times T_i$ be the space of type profiles, with generic element $t = (t_1, \ldots, t_n)$. We let $(T, \mathcal{T}, \mu)$ be a probability space, with each $t \in T$ representing a state of the world and μ being a probability measure.[8] Each voter i has a utility function $u_i \colon \mathcal{C} \times X \times T_i \to I\!R$, where $u_i(C, x, t_i)$ is continuous in x and measurable in t_i. Note that the voters' policy preferences are completely determined by their types, so that this is a setting of "private values," in the sense that each voter knows his or her preference when voting. Neverthless, the general form of the probability measure μ allows the model to handle arbitrary correlation structures in voters preferences, and so to admit cases where there could be substantial commonality in voters' preferences.

The payoff to voter i from outcome $(\bar{b}, \overline{x})$ is

$$u_i(\bar{b}, \overline{x}) \;=\; \sum_{C' \in W(\bar{b})} \frac{u_i(C, x_{C'}, t_i)}{\#W(\bar{b})}.$$

We say that voter i is *policy-oriented* if $u_i(C, x, t_i)$ is constant in its first argument.

Two prominent electoral models that are captured are:

- stochastic candidate bias: $T_i = I\!R^M$ and

$$u_i(C, x, t_i) \;=\; u_i(x) + t_{i,C}$$

- stochastic policy preference:

$$u_i(C, x, t_i) \;=\; u_i(x, t_i).$$

[8] As is clear in the proof, it is not necessary that μ be a probability measure - it could be a more general measure.

In the stochastic candidate bias model, the policy preferences of the voters are fixed, t is a $n \times M$ matrix, and the only unknown component in a voter i's utility is the term $t_{i,C}$ which embodies the incremental utility received by the voter if candidate C wins. It is commonly understood that this utility is derived from fixed characteristics of the candidates, or alternatively that the announced policies are not binding and by knowing which candidate is elected the voters can predict the policy that will ensue. In the stochastic policy preference model, a voter's utility depends not on the particular candidate who wins, but instead on that candidate's announced policy position.

Let us emphasize that our model is more general that the usual one, along a number of dimensions. For instance, in both of the cases above, it is typically assumed that there is a zero probability that voters are indifferent between the two candidates (provided the adopt distinct platforms in the stochastic policy preference model). We impose *no restriction* on the distribution of types. We also admit general forms of preferences for voters and candidates, in terms of how they value policy-candidate combinations.

A strategy for voter i is a behavioral strategy $\beta_i \colon \overline{X} \times T_i \to [0,1]$, where $\beta_i(x_A, x_B, t_i)$ represents the probability that voter i votes for candidate A given policy platforms x_A and x_B and type t_i. As usual, we require that β_i is measurable. We consider perfect Bayesian equilibrium of the electoral game in which voters use undominated* voting strategies.[9]

Theorem **5** *Consider the multi-candidate probabilistic voting model, where there are at least four voters or there are only two candidates. There exists a perfect Bayesian equilibrium in undominated* voting strategies. Moreover, there exists such an equilibrium that is Duvergerian, in which two arbitrary candidates A and B receive votes. If all voters are policy-oriented and candidates A and B are both win-motivated or both vote-motivated, then there exists such an equilibrium that is symmetric on $\{A, B\}$.*

If candidates are win-motivated or office-motivated, and if the probability of indifference between the candidates when they adopt distinct positions

[9]Note that because candidates are uninformed at the time that they adopt platforms, and because voters know their own preferences and thus have well-defined undominated* voting strategies (under the same definition as with complete information), the updating of beliefs does not play a role in the equilibrium analysis. Hence it is inconsequential whether we examine perfect Bayesian equilibrium or some stronger refinement.

is zero, then it can be shown that discontinuities in candidate payoffs are restricted to the diagonal. Thus, for an interesting subclass of models, discontinuities in payoffs take a very restricted form. In fact, candidate payoffs are fully continuous in the stochastic candidate bias model as long as voter types are continuously distributed. The aspect of Theorem 5 that admits discontinuities and requires the proof supplied here is that the distribution of voter preferences is arbitrary. This allows them to be indifferent in terms of how they vote, or to allow for correlation in their types.

Note also that the endogenous voting equilibrium approach can still be fruitful, even in the case where voters are (almost) never indifferent. If one considers limits of equilibria of probabilistic voting models as candidate uncertainty regarding voter preferences goes to zero; the limiting equilibria are themselves equilibria of the limiting complete information game. This can be proven either using Theorem 2 or Theorem 2 of Jackson, Simon, Swinkels and Zame (2002) (and noting that incentive compatibility is not an issue at any tie-breaking point). Thus, this offers another proof for existence of equilibrium in the two-candidate Downsian model: (i) examine a probabilistic voting version of the model, (ii) consider a sequence where the uncertainty in the stochastic candidate bias goes to zero, (iii) to select a corresponding sequence of equilibria in those models, and (iv) to take the limit of a convergent subsequence.

8 Public Randomization and Voting

We close with an extension of our results.

The importance of Duvergeian equilibrium in our analysis deals with convexity of the outcome correspondence. In particular, suppose that there are several possible candidates who could be elected given their platforms. If all voters are voting for just two candidates, then it must be that there are some indifferent voters. By varying the probabilities that they vote for either candidate, we can get any probability that one or the other wins. This allows us to convexify the outcome correspondence, which is important in establishing existence. When there are more than two candidates who can win in different equilibrium configurations of voters' behavior, it is possible that the outcome correspondence is not convex. This happens because voters are independently choosing their strategies, and full convexity over more than two candidates can require some correlation in how voters behave. Thus, to establish existence of equilibrium (in undominated strategies) where more than two candidates get votes, we need to have some means of correlating voters behavior. It is enough to have them observe a randomly generated public signal.

The definition of equilibrium with a publicly observed signal is again perfect Bayesian equilibrium.

THEOREM **6** *In the general probabilistic voting model, there exists a perfect Bayesian equilibrium with public randomization in undominated* voting strategies. Furthermore, if for all voters i, all distinct policies x and y, and all candidates A and B, $\mu(\{t \in T \mid u_i(A, x, t_i) = u_i(B, y, t_i)\}) = 0$, then there exists a perfect Bayesian equilibrium with public randomization in undominated voting strategies. If, in addition, for all voters i, all policies $x \in X$, and all candidates $A, B \in \mathcal{C}$, $\mu(\{t \in T \mid u_i(A, x, t_i) = u_i(B, x, t_i)\}) = 0$, then there exists a perfect Bayesian equilibrium in undominated voting strategies.*

The proof is a straightforward extension of the proofs of Theorems 3 and 5, and so we just discuss the critical step of showing convexity. Given the finite set of candidates, the convex hull of the set of possible equilibria of voters for any given policy choices of the candidates is defined by a finite set (say with cardinality K) of extreme points, with corresponding equilibria $\gamma^1, \ldots, \gamma^K$. A convex combination with weights $(\alpha^1, \ldots, \alpha^K)$ of these is obtained by having K different signals, each generated with these corresponding progabilities. Voters then play the equilibrium corresponding to the observed signal. Note that here since we are not starting with two a priori specified candidates A and B, we can start with equilibria in the closure of the set of undominated strategies. If types are almost never indifferent, then this equilibrium can be found so that voters play dominated strategies with probability zero.

Three features of Theorem 6 are worth noting. First, the general probabilistic voting model allows for a degenerate probability measure μ, and it therefore encompasses the multicandidate model with "deterministic" voting. Second, when the probability that voters are indifferent between distinct alternatives is zero, which precludes the deterministic model, we find equilibria in which voters eliminate undominated strategies in the standard sense. This is because subgames in which the undominated and undominated* strategies differ are probability zero from the perspective of candidates, allowing us to "patch up" our selection from equilibria in voting subgames. Third, our strongest assumption is satisfied in the stochastic candidate bias model when types are continuously distributed, but it is not satisfied in the stochastic policy preference model: there, all voters are indifferent between two candidates who adopt the same policy platform. Candidate payoffs are continuous in the candidate bias model when there are just two candidates, and we have noted that equilibrium existence is not an issue in that case. But with three or more candidates, it is not generally possible to take a continuous selection from equilibria in voting subgames, so the result of Theorem 6 is non-trivial even in that relatively tractable model.

9 Appendix A: Proofs

PROOF OF THEOREM 1: Let $E^{\nu}(\sigma(\bar{b}))$ denote the integral $\int \sigma(\bar{b})\nu(d\bar{b})$. The first part of Theorem 1 follows upon verifying the conditions of the main result in Simon and Zame (1990).[10] We claim that the payoff correspondence defined by

$$\{(u_A^{\psi}(\overline{x}), \ldots, u_M^{\psi}(\overline{x})) \mid \psi \text{ is a measurable section from } \Psi \}$$

is upper hemi-continuous with nonempty, compact, convex values. Nonemptiness is clear. Now take a sequence $\{\overline{x}^r\}$ in $\overline{X}$ converging to $\overline{x}$ and a sequence ψ^r of measurable selections from Ψ such that $\{u_C^{\psi^r}(\overline{x}^r)\}$ converges for all $C \in \mathcal{C}$. Since $\Delta(\mathcal{C}^n)$ is compact, there is a subsequence of $\{\psi^r(\overline{x}^r)\}$, which we continue to index by r for simplicity, that converges to some distribution $\nu \in \Delta(\mathcal{C}^n)$. Since Ψ has closed graph, it follows that $\nu \in \Psi(\overline{x})$. Letting ψ be any measurable selection of Ψ such that $\psi(\overline{x}) = \nu$, we have $u_C^{\psi}(\overline{x}) = \lim u_C^{\psi^r}(\overline{x}^r)$ for all $C \in \mathcal{C}$, and $(u_A^{\psi}(\overline{x}), \ldots, u_M^{\psi}(\overline{x})) \in \Psi(\overline{x})$, establishing closed graph. To establish convex values, take any $\overline{x}$, any $\alpha \in (0, 1)$, and any measurable selections ψ and ψ' from Ψ. Let $\nu, \nu' \in \Psi(\overline{x})$ satisfy $\nu = \psi(\overline{x})$ and $\nu' = \psi'(\overline{x})$. For all $C \in \mathcal{C}$, linearity of u_C in s implies

$$\alpha u_C^{\psi}(\overline{x}) + (1-\alpha) u_C^{\psi'}(\overline{x}) \quad = \quad u_C(\alpha E^{\nu}(\sigma(\bar{b})) + (1-\alpha) E^{\nu'}(\sigma(\bar{b})), \overline{x}).$$

[10]Simon and Zame (1990) work with a correspondence of utilities rather than outcomes. Theorem 1 in Jackson, Simon, Swinkels and Zame (2002) works with outcome correspondences and directly implies the first statement in Theorem 1.

By the assumption that $\sigma\Psi(\overline{x})$ is convex, we can choose $\hat{\nu} \in \Psi(\overline{x})$ independent of C such that

$$E^{\hat{\nu}}(\sigma(\overline{b})) \;=\; \alpha E^{\nu}(\sigma(\overline{b})) + (1-\alpha)E^{\nu'}(\sigma(\overline{b})).$$

Let $\hat{\psi}$ be a measurable selection of Ψ satisfying $\hat{\psi}(\overline{x}) = \hat{\nu}$. Then

$$\alpha u_C^{\psi}(\overline{x}) + (1-\alpha)u_C^{\psi'}(\overline{x}) \;=\; u_C^{\hat{\psi}}(\overline{x}),$$

which establishes convex values, as required. The second part then follows from Theorem 2 in Jackson, Simon, Swinkels, and Zame (2002, remark (ii)). ∎

PROOF OF THEOREM 3: Pick two candidates A and B from $\mathcal{C}$. Consider a restricted game where voters can only vote for one of these two candidates, and we model only these two candidates' strategies. We prove existence of an equilibrium in this modified game. Note that if there are more candidates, and at least four voters, then the following is a subgame perfect equilibrium in undominated* voting strategies of the original game: have voters follow the strategies from the restricted game independent of the policies of the remaining candidates, and have the remaining candidates play arbitrary strategies. Note that the role of having at least four voters is that any deviation by a single voter cannot result in the election of any other candidate. Thus, we need only prove that there exists an equilibrium in the two candidate game.

Construct a voting correspondence Ψ as follows. Let $\Psi(x_A, x_B)$ consist of all distributions $\nu \in \Delta(\mathcal{C}^n)$ such that for all voters i, there exists $\gamma_i \in [0,1]$ satisfying

(1) for all $\overline{b} \in \mathcal{C}^n$,

$$\nu(\overline{b}) \;=\; \int_T \left(\prod_{i:b_i=A} \gamma_i\right)\left(\prod_{i:b_i=B} (1-\gamma_i)\right)\mu(dt)$$

(2) for all $i \in \mathcal{V}$,

$$\gamma_i \;=\; 1 \quad \text{if } u_i(A,x_A) > u_i(B,x_B)$$
$$\gamma_i \;=\; 0 \quad \text{if } u_i(A,x_A) < u_i(B,x_B).$$

Let $\gamma = (\gamma_1, \ldots, \gamma_n)$. Here, we interpret γ_i as the probability that voter i votes for candidate A as a function of the voter's type, where we require the voter to put probability zero on dominated votes. This then determines the probability that any particular ballot vector is realized.

Clearly, Ψ has nonempty values. We claim that, for all $\overline{x} \in \overline{X}$, $\Psi(\overline{x})$ is connected. Given ν generated by γ and ν' generated by γ', let ν^α be the distribution generated by $\alpha\gamma + (1-\alpha)\gamma'$. Then the path swept out by ν^α as α varies over $[0,1]$ connects ν and ν', establishing the claim. Note that $E^\nu(\sigma(\bar{b}))$ is continuous in ν, and therefore $\sigma\Psi(\overline{x})$, as the image of a connected set under a continuous function, is connected. By our dimensionality assumption, the range of $E^\nu(\sigma(\bar{b}))$ is one-dimensional, and it follows that $\sigma\Psi(\overline{x})$ is indeed convex. Therefore, $\sigma\Psi$ has convex values.

To verify that $\sigma\Psi$ has closed graph, take a sequence (x_A^k, x_B^k, s^k) such that $(x_A^k, x_B^k, s^k) \to (x_A, x_B, s)$ and, for all k, $s^k \in \sigma\Psi(x_A^k, x_B^k)$. By the latter, for each k, there exists $\nu^k \in \Psi(x_A^k, x_B^k)$ such that $s^k = E^{\nu^k}(\sigma(\bar{b}))$. Thus, for each k, there exists γ^k satisfying (1) and (2) above with x_A^k and x_B^k in condition (2). By compactness, there is a convergent subsequence of $\{\gamma^k\}$, still indexed by k for simplicity. Let γ denote the limit of this subsequence, and let ν be defined from γ as in condition (1), so that $s = E^\nu(\sigma(\bar{b}))$. Then continuity of u_i ensures that condition (2) also holds, and therefore $s \in \Psi(x_A, x_B)$. Thus, $\sigma\Psi$ has closed graph.

By Theorem 1, there is an endogenous voting equilibrium (ξ, ψ), where ψ is a selection from Ψ. Define the behavioral strategy β_i for voter i as follows. For all $x_A, x_B \in X$, since ψ selects from Ψ, there exists γ satisfying (1) and (2) with $\nu = \psi(x_A, x_B)$. Then let $\beta_i(x_A, x_B) = \gamma_i)$, completing the equilibrium construction and proving the first part of the theorem.

For the second part of the theorem, note that if voters are policy-oriented and A and B are both win-motivated or both vote-motivated, then the electoral game $(\overline{X}, \Psi, u)$ is symmetric on $\{A, B\}$, and Theorem 1 yields an endogenous voting equilibrium that is symmetric on $\{A, B\}$. We then use the associated selection to define symmetric voting strategies, as required. ∎

We now turn to the proof of Theorem 4. First we present a useful lemma. Our equilibrium analysis uses the following dominance relation on the candidates' strategies. We say x *stage dominates* y, written $x \, SD \, y$, if, for every symmetric, undominated β, we have $u_A^\beta(x, z) \geq u_A^\beta(y, z)$ for all $z \in X$, with strict inequality for at least one $z \in X$ (where the first entry in $u_A^\beta(x,z)$ is A's chosen policy). Note that, given undominated β, xPy implies that a majority of voters must vote for the candidate, say A, with position x, implying $u_A^\beta(x,y) = 1$ and $u_B^\beta(x,y) = -1$ (noting that in the Downsian setting the election is a zero sum game and assigning a value of 1 to winning the election). The next result establishes a tight connection between deep covering and stage dominance of the electoral model.

LEMMA **1** *In the two-candidate Downsian model, $x\,DC\,y$ implies $x\,SD\,y$. Assuming n is odd, the converse holds as well.*

PROOF OF LEMMA 1: Assume $x\,DC\,y$, and take any $z \in X$. If xPz, then $u_A^\beta(x,z) = 1$, and the weak inequality in the definition of stage dominance is fulfilled. If zRx, then, by deep covering, zPy. This implies $u_A^\beta(y,z) = -1$, fulfilling the weak inequality in stage dominance. Finally, when $z = x$, we have $u_A^\beta(x,z) = 0$ by symmetry of β, and, by xPy, we have $u_A^\beta(y,z) = -1$, fulfilling the strict inequality in stage dominance. Therefore, $x\,SD\,y$. Now assume n is odd and $x\,SD\,y$, and take any $z \in R(x)$. Suppose $z \notin P(y)$, i.e., yRz. Since n is odd, we have

$$|\{i \mid u_i(z) \geq u_i(x)\}| > \frac{n}{2} \quad \text{and} \quad |\{i \mid u_i(y) \geq u_i(z)\}| > \frac{n}{2}.$$

Therefore, there exists a symmetric, undominated β such that $u_A^\beta(x,z) = -1$ and $u_A^\beta(y,z) = 1$, a contradiction. Therefore, $z \in P(y)$, and we conclude that $x\,DC\,y$. ∎

An immediate corollary, using the above observations on the deep uncovered set, is that when n is odd the set of stage undominated policy positions is non-empty and compact. When n is even, it need not be the case that stage dominance implies deep covering: in this case, we may have $x\,SD\,y$ yet xIz and yIz for some z, if no voters are indifferent between x and z and none indifferent between y and z – so majority indifference is the result of an equal split among the voters, with all voters having strict preferences. In this case, the latitude in specifying the behavior of indifferent voters does not play a role, and no contradiction need arise.

Proof of Theorem 4: Let $(\hat{\xi}_A, \hat{\xi}_B, \beta)$ be a subgame perfect equilibrium in symmetric, undominated voting strategies, and suppose that the support of $\hat{\xi}_A$ is not contained in DC. Letting V denote the set of positions that are deeply covered, this means $\hat{\xi}_A(V) > 0$. Let $\xi_A = \xi_B = \hat{\xi}_A$. Since the game between the candidates induced by β is symmetric and zero-sum, (ξ_A, ξ_B, β) is a subgame perfect equilibrium in symmetric, undominated voting strategies. For each $y \in V$, there exists $z \in X$ such that $z \, DC \, y$. And since the deep covering relation DC is open, there exists an open set $G_y \subseteq X$ such that $y \in G_y$ and $z \, DC \, G_y$. Giving V the relative topology, then $\{G_y \cap V \mid y \in V\}$ is an open cover of V. Since X is second countable, V is as well, and it therefore possesses the Lindelöf property: so there is a countable open subcover given by, say, $\{H_k\}$. Since $\xi_A(V) > 0$, it follows that $\xi_A(H_k) > 0$ for some k. Let $z_k \in X$ satisfy $z_k \, DC \, H_k$, and transform ξ_A to ξ'_A by "replacing" any policy in H_k with z_k. Specifically, define the mapping $\phi \colon X \to X$ as follows:

$$\phi(x) \;=\; \begin{cases} z_k & \text{if } x \in H_k \\ x & \text{else,} \end{cases}$$

for all x. Define $\xi'_A = \xi_A \circ \phi^{-1}$, where $(\xi_A \circ \phi^{-1})(Y) = \xi_A(\phi^{-1}(Y))$, for all Borel measurable $Y \subseteq X$. Then

$$\begin{aligned} & U_A^\beta(\xi'_A, \xi_B, \beta) \\ &= \int u_A^\beta(\overline{x}) \, (\xi'_A \times \xi_B)(d\overline{x}) \\ &= \int u_A^\beta(\phi(x_A), x_B) \, \xi(d\overline{x}) \\ &= \int_{H_k \times H_k} u_A^\beta(\phi(x_A), x_B) \, \xi(d\overline{x}) \\ &\quad + \int_{(X \times X) \setminus (H_k \times H_k))} u_A^\beta(\phi(x_A), x_B) \, \xi(d\overline{x}). \end{aligned}$$

Note that $z_k = \phi(x_A) \, DC \, H_k$ for all $x_A \in H_k$, so $u_A^\beta(\phi(x_A), x_B) = 1$ for all $x_A, x_B \in H_k$, which implies

$$\int_{H_k \times H_k} u_A^\beta(\phi(x_A), x_B) \, \xi(d\overline{x}) \;=\; \xi_A(H_k) \xi_B(H_k) \;>\; 0.$$

In contrast, symmetry implies

$$\int_{H_k \times H_k} u_A^\beta(x_A, x_B) \, \xi(d\overline{x}) \;=\; 0.$$

From Lemma 1, we have $z_k = \phi(x_A)\, SD\, H_k$ for all $x_A \in H_k$, which implies

$$\int_{(X\times X)\setminus(H_k\times H_k))} u_A^\beta(\phi(x_A), x_B)\,\xi(d\overline{x}) \;\geq\; \int_{(X\times X)\setminus(H_k\times H_k))} u_A^\beta(x_A, x_B)\,\xi(d\overline{x}).$$

Combining these observations yields

$$\begin{aligned} U_A^\beta(\xi'_A, \xi_B, \beta) &> \int_{H_k\times H_k} u_A^\beta(x_A, x_B)\,\xi(d\overline{x}) \\ &\quad + \int_{(X\times X)\setminus(H_k\times H_k)} u_A^\beta(x_A, x_B)\,\xi(d\overline{x}) \\ &= U_A^\beta(\xi_A, \xi_B, \beta), \end{aligned}$$

contradicting the assumption that (ξ_A, ξ_B, β) is an equilibrium. ∎

Proof of Theorem 5: Just as in Theorem 3, we pick two candidates A and B from $\mathcal{C}$. Consider a restricted game where voters can only vote for one of these two candidates, and we model only these two candidates' strategies. We prove existence of an equilibrium in this modified game. Note that if there are more candidates, and at least four voters, then the following is a perfect Bayesian equilibrium in undominated* voting strategies of the original game: have voters follow the strategies from the restricted game independent of the policies of the remaining candidates, and have the remaining candidates play arbitrary strategies. Note that the role of having at least four voters is that any deviation by a single voter cannot result in the election of any other candidate. Thus, we need only prove that there exists an equilibrium in the two candidate game.

Construct a voting correspondence Ψ as follows. Let $\Psi(x_A, x_B)$ consist of all distributions $\nu \in \Delta(\mathcal{C}^n)$ such that for all voters i, there exists a measurable mapping $\gamma_i\colon T_i \to [0,1]$ satisfying

(1) for all $\overline{b} \in \mathcal{C}^n$,

$$\nu(\overline{b}) \;=\; \int_T \left(\prod_{i:b_i=A} \gamma_i(t_i)\right)\left(\prod_{i:b_i=B} (1-\gamma_i(t_i))\right)\mu(dt)$$

(2) for all $i \in \mathcal{V}$ and all $t_i \in T_i$,

$$\begin{aligned} \gamma_i(t_i) &= 1 \;\text{ if } u_i(A, x_A, t_i) > u_i(B, x_B, t_i) \\ &= 0 \;\text{ if } u_i(A, x_A, t_i) < u_i(B, x_B, t_i). \end{aligned}$$

Let $\gamma = (\gamma_1, \ldots, \gamma_n)$. Here, we interpret γ_i as the probability that voter i votes for candidate A as a function of the voter's type, where we require the voter to put probability zero on dominated votes. This then determines the probability that any particular ballot vector is realized.

Because u_i is measurable, Ψ has nonempty values. We claim that, for all $\overline{x} \in \overline{X}$, $\Psi(\overline{x})$ is connected. Given ν generated by γ and ν' generated by γ', let ν^α be the distribution generated by $\alpha\gamma + (1-\alpha)\gamma'$. Then the path swept out by ν^α as α varies over $[0,1]$ connects ν and ν', establishing the claim. Note that $E^\nu(\sigma(\bar{b}))$ is continuous in ν, and therefore $\sigma\Psi(\overline{x})$, as the image of a connected set under a continuous function, is connected. By our dimensionality assumption, the range of $E^\nu(\sigma(\bar{b}))$ is one-dimensional, and it follows that $\sigma\Psi(\overline{x})$ is indeed convex. Therefore, $\sigma\Psi$ has convex values.

To verify that $\sigma\Psi$ has closed graph, take a sequence (x_A^k, x_B^k, s^k) such that $(x_A^k, x_B^k, s^k) \to (x_A, x_B, s)$ and, for all k, $s^k \in \sigma\Psi(x_A^k, x_B^k)$. By the latter, for each k, there exists $\nu^k \in \Psi(x_A^k, x_B^k)$ such that $s^k = E^{\nu^k}(\sigma(\bar{b}))$. Thus, for each k, there exists γ^k satisfying (1) and (2) above with x_A^k and x_B^k in condition (2). We must construct ν and γ satisfying (1) and (2) for the limiting platforms x_A and x_B and such that $s = E^\nu(\sigma(\bar{b}))$.

For each i, let $T_i^>$ denote the set of types for which i prefers candidate A with platform x_A to B with platform x_B, let $T_i^<$ denote the set of types for which i prefers B, and let $T^=(i)$ be the set of types such that i is indifferent:

$$\begin{aligned} T_i^> &= \{t_i \in T_i \mid u_i(A, x_A, t_i) > u_i(B, x_B, t_i)\} \\ T_i^< &= \{t_i \in T_i \mid u_i(A, x_A, t_i) < u_i(B, x_B, t_i)\} \\ T_i^= &= \{t_i \in T_i \mid u_i(A, x_A, t_i) = u_i(B, x_B, t_i)\}, \end{aligned}$$

and let $T^= = \bigcup_{i \in \mathsf{V}} (T_i^= \times T_{-i})$ denote the set of type profiles at which some voter is indifferent. Take any $t \in T_i^>$. By continuity of u_i, voter i prefers candidate A for sufficiently high k, and therefore $\gamma_i^k(t_i) = 1$; similarly, if $t_i \in T_i^<$, then $\gamma_i^k(t_i) = 0$ for sufficiently high k. Given $\alpha \in [0,1]^n$, define $\gamma_i^\alpha \colon T_i \to [0,1]$ by

$$\gamma_i^\alpha(t_i) = \begin{cases} 1 & \text{if } t_i \in T_i^> \\ \alpha_i & \text{if } t \in T_i^= \\ 0 & \text{if } t \in T_i^<, \end{cases}$$

which is clearly measurable. Letting $\gamma^\alpha = (\gamma_1^\alpha, \ldots, \gamma_n^\alpha)$, we see that γ^α fulfills condition (2). Let ν^α be the distribution determined by γ^α through condition (1).

We have left to establish existence of $\alpha \in [0,1]^n$ such that $E^{\nu^\alpha}(\sigma(\bar{b})) = s$. For convenience, define

$$\Gamma^k(\bar{b}|t) \;=\; \left(\prod_{i:b_i=A} \gamma_i^k(t_i)\right)\left(\prod_{i:b_i=B}(1-\gamma_i^k(t_i))\right),$$

and define $\Gamma^\alpha(t)$ similarly using γ^α. Note that

$$\begin{aligned} E^{\nu^k}(\sigma(\bar{b})) &= \int_T \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^k(\bar{b}|t)\mu(dt) \\ E^{\nu^\alpha}(\sigma(\bar{b})) &= \int_T \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^\alpha(\bar{b}|t)\mu(dt). \end{aligned}$$

We have argued that each γ_i^k converges pointwise to γ_i^α on $T \setminus T^=$, and therefore $\Gamma^k(\cdot|\bar{b})$ also converges pointwise to $\Gamma^\alpha(\cdot|\bar{b})$ on $T \setminus T^=$. As a consequence, we have

$$\int_{T\setminus T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^\alpha(\bar{b}|t)\mu(dt) \;=\; \lim_{k\to\infty} \int_{T\setminus T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^k(\bar{b}|t)\mu(dt).$$

If $\mu(T^=) = 0$, then the above equality yields $E^{\nu^\alpha}(\sigma(\bar{b})) = s$, as desired.

Otherwise, if $\mu(T^=) > 0$, define

$$\theta \;=\; \lim_{k\to\infty} \frac{1}{\mu(T^=)} \int_{T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^k(b|t)\mu(dt).$$

By our dimensionality and monotonicity assumptions, we may assume that $\bar{b} \geq_A \bar{b}'$ implies $\sigma(\bar{b}) \geq \sigma(\bar{b}')$. We therefore have

$$\frac{1}{\mu(T^=)} \int_{T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^0(b|t)\mu(dt) \;\leq\; \theta \;\leq\; \frac{1}{\mu(T^=)} \int_{T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^1(b|t)\mu(dt).$$

Thus, by continuity, we may choose $\alpha \in [0,1]$ such that

$$\theta \;=\; \frac{1}{\mu(T^=)} \int_{T^=} \sum_{\mathsf{C}^n} \sigma(\bar{b})\Gamma^\alpha(b|t)\mu(dt),$$

yielding α such that $E^{\nu^\alpha}(\sigma(\bar{b})) = s$, completing the proof of closed graph.

By Theorem 1, there is an endogenous voting equilibrium (ξ, ψ), where ψ is a selection from Ψ. Define the behavioral strategy β_i for voter i as follows. For all $x_A, x_B \in X$, since ψ selects from Ψ, there exists γ satisfying (1) and (2) with $\nu = \psi(x_A, x_B)$. Then let $\beta_i(x_A, x_B, t_i) = \gamma_i(t_i)$, completing the equilibrium construction and proving the first part of the theorem.

For the second part of the theorem, note that if voters are policy-oriented and A and B are both win-motivated or both vote-motivated, then the electoral game $(\overline{X}, \Psi, u)$ is symmetric, and Theorem 1 yields a symmetric endogenous voting equilibrium. We then use the associated selection to define symmetric voting strategies, as required. ∎

10 Appendix B: An Asymmetric Voting Example

It is relatively easy to construct examples that violate Theorem 4, once the restriction of symmetric voting is removed. For example, suppose that all voters have the same ideal point and vote for candidate A if A proposes that policy, and otherwise specify any undominated votes; let candidate A choose the voters' ideal point, and let candidate B choose any policy position. This is a subgame perfect equilibrium in undominated voting strategies, but it exhibits two features that make it uninteresting. First, all voters vote for A, even if both candidates adopt the voters' ideal point, an extreme violation of symmetric voting. Second, the position chosen by candidate B, which may be covered, does not arise as a policy outcome, since A wins the election with probability one.

The example in Table 1 shows that, even if we maintain the assumption that voters flip a fair coin to decide their vote in case both candidates adopt the same position, it is possible that one candidate adopts a covered policy position. In fact, in the example, that candidate wins with positive probability after adopting that position.

Here, let rows denote the strategies of candidate A, columns the strategies of B, and let entries denote A's payoffs: "0" denotes majority indiffer-

	a	b	c	d	e	f	g
a	0	0	0	-1	-1	1	$0^{(-1)}$
b	0	0	0	-1	$0^{(-1)}$	-1	1
c	0	0	0	1	1	$0^{(-1)}$	-1
d	1	1	-1	0	1	-1	-1
e	1	$0^{(1)}$	-1	-1	0	-1	-1
f	-1	1	$0^{(1)}$	1	1	$0^{(1)}$	$0^{(1)}$
g	$0^{(1)}$	-1	1	1	1	$0^{(1)}$	$0^{(1)}$

Table 1: Covered equilibrium platform

ence, and a superscript indicates the resolutions of indifferent voters. For example, "$0^{(1)}$" indicates that a sufficient number of indifferent voters vote for A with probability one that A wins the election with certainty. Specify mixed strategies for the candidates so that $\xi_A = (0,0,0,0,\frac{1}{3},\frac{1}{3},\frac{1}{3})$ and $\xi_B = (\frac{1}{3},\frac{1}{3},\frac{1}{3},0,0,0,0)$. It is straightforward to check that these strategies, with the voters' strategies implicit in the candidates' payoffs, form a subgame perfect equilibrium (and we may assume that voters use undominated strategies). Note, however, that the voters' strategies do not treat the candidates symmetrically: when A chooses g and B chooses f, voters resolve indifference in favor of A, and they also do so when A chooses f and B chooses g.

In this example, candidate A chooses policy e with probability one third, despite the fact that it is covered by d: driving this is the fact that the voter votes for A with probability one when A chooses e and B chooses b, giving A an expected payoff of 1/3 from choosing e (as with f or g). It then becomes critical to limit B's payoff to $-1/3$ if that candidate deviates to e, f, or g, necessitating the asymmetric voting strategies employed in the example.

Note that A's expected payoff upon choosing e, as a function of B's choice of a, b, or c, is actually identical to the payoff the candidate would receive upon choosing d. In fact, letting A choose d with probability one third instead of e, we have an "equivalent" equilibrium in which the candidates use only uncovered strategies. Theorem 4 can be generally extended in this way to allow for asymmetric voting, when appropriately restated: for every subgame perfect equilibrium in undominated voting strategies, say (ξ, β), there exists another such equilibrium, say (ξ', β), such that ξ'_A and ξ'_B have support on the deep uncovered set and the distribution of each candidate's payoffs are equivalent to that under (ξ, β).

References

[1]

[2] J. Banks, J. Duggan, and M. Le Breton (2002) "Bounds for Mixed Strategy Equilibria and the Spatial Model of Elections," *Journal of Economic Theory,* 103: 88-105.

[3] J. Banks, J. Duggan, and M. Le Breton (2003) "Social Choice and Electoral Competition in the General Spatial Model," mimeo.

[4] D. Black. 1958. *The Theory of Committees and Elections.* Cambridge: Cambridge University Press.

[5] G. Bordes, M. Le Breton, and M. Salles (1992) Gillies and Miller's Subrelations of a Relation over an Infinite Set of Alternatives: General Results and Applications to Voting Games, *Math. Op. Res.* 17, 509-518.

[6] P. Dasgupta and E. Maskin (1986) "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," *Review of Economic Studies,* 53: 1-26.

[7] A. Downs. 1957. *An Economic Theory of Democracy.* New York: Harper and Row.

[8] J. Duggan (2005) "Equilibrium Existence in Discontinuous Zero-sum Games with Applications to Spatial Models of Elections," mimeo.

[9] J. Duggan (2006) "Uncovered Sets," mimeo.

[10] B. Dutta and J.-F. Laslier (1999) "Comparison Functions and Choice Correspondences," *Soc. Choice Welfare* 16, 513-532.

[11] P. Fishburn (1977) "Condorcet Social Choice Functions," *SIAM J. Applied Math.* 33, 469-489.

[12] M.O. Jackson, L. Simon, J. Swinkels, and W. Zame (2002) "Communication and Equilibrium in Discontinuous Games of Incomplete Information," *Econometrica,* 70, 1711-1740.

[13] M.O. Jackson and J. Swinkels (2005) "Existence of Equilibria in Single and Double Auctions," *Econometrica*, 73(1), 93-139.

[14] G. Kramer (1978) "Existence of Electoral Equilibrium, in "Game Theory and Political Science" (P. Ordeshook, ed.), New York University Press, New York.

[15] G. Laffond, J.-F. Laslier, and M. Le Breton (1993) "The Bipartisan Set of a Tournament Game," *Games Econ. Behav.* 5, 182-201.

[16] R. McKelvey (1986) "Covering, Dominance, and Institution-free Properties of Social Choice," *Amer. J. Pol. Sci.* 30, 283-314.

[17] N. Miller (1980) "A New Solution Set for Tournaments and Majority Voting: Further Graph-theoretical Approaches to the Theory of Voting," *Amer. J. Pol. Sci.* 24, 68-96.

[18] J. Nash (1950) "Equilibrium Points in n-person Games," *Proceedings of the National Academy of Sciences,* 36, 48-49.

[19] C. Plott (1967) "A Notion of Equilibrium and Its Possibility Under Majority Rule," *Amer. Econ. Rev.* 57, 787-806.

[20] P. Reny (1999) "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games," *Econometrica,* 67: 1029-1056.

[21] L. Simon and W. Zame (1990) "Discontinuous Games and Endogenous Sharing Rules," *Econometrica,* 58, 861-872.