A Survey of Equilibrium Analysis in Spatial Models of Elections*

John Duggan
Department of Political Science
and Department of Economics
University of Rochester

July 21, 2005

1 Introduction

Elections, as the institution through which citizens choose their political agents, are at the core of representative democracy. It is therefore appropriate that they occupy a central position in the study of democratic politics. The formal analysis of elections traces back to the work of Hotelling (1929), Downs (1957), and Black (1958), who apply mathematical methods to understand the equilibrium outcomes of elections. This work, and the literature stemming from it, has focused mainly on the positional aspects of electoral campaigns, where we conceptualize candidates as adopting positions in a "space" of possible policies prior to an election. We maintain this focus by considering the main results for the canonical model of elections, in which candidates simultaneously adopt policy platforms and the winner is committed to the platform on which he or she ran. These models abstract from much of the structural detail of elections, including party primary elections, campaign finance and advertising, the role of interest groups, etc. Nevertheless, in order to achieve a deep understanding of elections in their full complexity, it seems that we must address the equilibrium effects of position-taking by candidates in elections.

In this article, I will cover known foundational results on spatial models of elections, taking up issues of equilibrium existence, the distance (or lack thereof) between the equilibrium policy positions of the candidates, and the characterization of equilibria in terms of social choice concepts such as the majority core

*I thank David Epstein, Mark Fey, Cesar Martinelli, and Martin Osborne for helpful feedback.

and the utilitarian social welfare function. The article is structured primarily by assumptions on voting behavior. I first consider results for the case of deterministic voting, which I refer to as the "Downsian model," where candidates can essentially predict the votes of voters following policy choices. I then consider two models of probabilistic voting, where voting behavior is modelled as a random variable from the perspective of the candidates. Within each section, I consider the most common objective functions used to model the electoral incentives of different types of candidates, including candidates who seek only to win the election, candidates who seek only to maximize their vote totals, and candidates who seek the best policy outcome from the election.

Of several themes in the article, most prominent will be difficulties in obtaining existence of equilibrium, especially when the policy space is multidimensional. While the median voter theorem establishes existence of an equilibrium in the unidimensional Downsian model, models of probabilistic voting are commonly thought to mitigate existence problems in multiple dimensions. We will see that this is true to an extent, but that probabilistic voting can actually introduce existence problems in the unidimensional model.

As a point of reference for the equilibrium existence issue, early articles by Debreu (1952), Fan (1952), and Glicksberg (1952) give useful sufficient conditions for existence of equilibrium in the games we analyze.[1] Their existence result, which we will refer to as the "DFG theorem," first assumes each player's set of strategies is a subset of $\Re^n$ that is non-empty, compact (so it is described by a well-defined boundary in $\Re^n$), and convex (so a player may move from one strategy toward any other with no constraints). These regularity assumptions that are easily satisfied in most models. Second, and key to our analysis, DFG assumes that the objective function of each player is:

- jointly continuous in the strategies of all players (so small changes in the strategies of the players lead to small changes in payoffs)
- quasi-concave in that player's own strategy, given any strategies for the other players (so any move toward a better strategy increases a player's payoff).

These continuity and convexity conditions are violated in a range of electoral models. The possibility of discontinuities is well-known, and it is often blamed for existence problems. Probabilistic voting models smooth the objective functions of the candidates, preventing such discontinuities, but equilibrium existence can still be problematic due to convexity problems. Thus, while the issue of convexity may receive less attention than continuity, it is equally critical in obtaining existence of equilibria.

[1] Nash (1950) proves existence of mixed strategy equilibrium for finite games. Since our games involve convex (and therefore infinite) policy spaces, his result does not apply here.

Quasi-concavity of the candidates' objective functions is only an issue, however, if we seek equilibria in *pure* strategies. More generally, we may allow the candidates to use *mixed strategies,* which formalize the idea that neither candidate can precisely predict the campaign promises of the other. In this case, the results of DFG can be used to drop quasi-concavity: continuity of the objective functions is sufficient for existence of mixed strategy equilibria. In continuous models, therefore, this offers one solution to the existence problem. In discontinuous models, which include most of those we cover, it is often possible to appeal to even more general sets of sufficient conditions to obtain mixed strategy equilibria. In resorting to mixed strategies to solve the existence problem, we arrive at a conclusion that contrasts Riker's (1980) claim of the inherent instability of democratic politics: rather than taking the absence of pure strategy equilibria as evidence of instability, we acknowledge that the positional aspects of elections give the candidates an incentive to be unpredictable, as is the case for players in many other strategically complex games. This element of indeterminacy does not, however, preclude a scientific approach to the analysis of elections, as it is still possible to make statistical predictions, to give bounds on the possible policy positions, and to perform comparative statics.

For further background on the electoral modelling literature, there are a number of surveys, such as Wittman (1990), Coughlin (1990, 1992), and Austen-Smith and Banks (2004). I do not touch on a number of interesting issues in electoral modelling, such as multiple (three or more) candidates, entry and exit of candidates, informational aspects of campaigns. See Calvert (1986), Shepsle (1991), and Osborne (1995) for surveys of much of that work.

2 The Electoral Framework

We will focus on the spatial model of politics, as elaborated by Davis, Hinich, and Ordeshook (1970), Ordeshook (1986), and Austen-Smith and Banks (1999). Here, we assume that the policy space X is a subset of Euclidean space of some finite dimension, d. Thus, a policy is a vector $x = (x_1, \ldots, x_d)$, where x_k may denote the amount of spending on some project or a position on some issue, suitably quantified. We assume that X is non-empty, compact, and convex. We consider an election with just two political candidates, A and B (sometimes interpreted as parties), and we analyze an abstract model of campaigns: we assume that the candidates simultaneously announce policy positions x_A and x_B in the space X, and we assume that the winning candidate is committed to his or her campaign promise. Following these announcements, a finite, odd number n of voters, denoted $i = 1, 2, \ldots, n$, cast their ballots $b_i \in \{1, 0\}$, where $b_i = 1$ denotes a vote for candidate A and $b_i = 0$ denotes a vote for B, the winner being the candidate with the most votes. We do not allow abstention by voters.

For now, we model the objectives of the candidates and the behavior of voters at a general level. Let $U_A(x_A, x_B, b_1, \ldots, b_n)$ denote candidate A's utility when the candidates take positions x_A and x_B and the vector of ballots is $(b_1, \ldots, b_n)$, and define the notation $U_B(x_A, x_B, b_1, \ldots, b_n)$ similarly. At this level of generality, we allow for the possibility that candidates seek only to win the election, or that they seek only the most favorable policies possible, or a mix of these ambitions. Before proceeding to describe these objective functions precisely, let

$$\Phi(b_1, \ldots, b_n) \quad = \quad \begin{cases} 1 & \text{if } \sum_i b_i > \frac{n}{2} \\ 0 & \text{if } \sum_i b_i < \frac{n}{2} \end{cases}$$

indicate whether a majority of voters have voted for candidate A or B.

The first main approach to modelling candidate objectives is to view the candidates as primarily concerned with their electoral prospects. This is reflected in our first objective function, which dictates that a candidate cares only about whether he or she garners a majority of votes. In other words, only the sign, rather than the magnitude, of the margin of victory matters.

Win motivation. The candidates receive utility equal to one from winning, zero otherwise, so that

$$U_A(x_A, x_B, b_1, \ldots, b_n) \quad = \quad \Phi(b_1, \ldots, b_n),$$

with candidate B's utility equal to one minus the above quantity. According to the second version of office motivation, the candidates do indeed care about the margin of victory.

The second objective function also captures the idea that candidates care primarily about electoral success, but now measured in the number of votes for the candidate.

Vote motivation. The candidates' utilities take the simple linear form

$$U_A(x_A, x_B, b_1, \ldots, b_n) \quad = \quad \sum_i b_i,$$

with candidate B's utility being $n - \sum_i b_i$. Note that, since we rule out abstention by voters, vote motivation is equivalent by a positive affine transformation to plurality motivation: the margin of victory for candidate A, for example, is just $(2 \sum b_i) - n$.

By the term *office motivation,* we refer to the situation in which either both candidates are win-motivated or both are vote-motivated. The second main approach to modelling candidate objectives is to assume that candidates care only about the policy outcome of the election.

Policy motivation. Assume the candidates have policy preferences represented by strictly concave, differentiable utility functions u_A and u_B. In words, the graphs of these utility functions are smooth and "dome-shaped." In particular, utilities are quasi-concave. Under these assumptions, each candidate has an *ideal policy,* which yields a strictly higher utility than all other policies, and we denote these as $\tilde{x}_A$ and $\tilde{x}_B$. We assume these ideal policies are distinct, and in the special case of a unidimensional policy space, we assume $\tilde{x}_A < \tilde{x}_B$ without loss of generality. Then candidate A's utility from policy positions x_A and x_B and ballots $(b_1, \ldots, b_n)$ is

$$\begin{aligned} & U_A(x_A, x_B, b_1, \ldots, b_n) \\ = & \quad \Phi(b_1, \ldots, b_n) u_A(x_A) + (1 - \Phi(b_1, \ldots, b_n)) u_A(x_B), \end{aligned}$$

and likewise for candidate B.

Though I focus on pure office and policy motivation in this article, a third approach which combines the above two has also been considered in the literature.

Mixed motivation. Assume that candidates have policy preferences as above, and that the winner of the election receives a reward $w > 0$. Then candidate A's utility is

$$\begin{aligned} & U_A(x_A, x_B, b_1, \ldots, b_n) \\ = & \quad \Phi(b_1, \ldots, b_n)(u_A(x_A) + w) + (1 - \Phi(b_1, \ldots, b_n)) u_A(x_B), \end{aligned}$$

and likewise for B.

In order to model voting behavior, let $P_i(x_A, x_B)$ denote the probability that voter i votes for candidate A, i.e., casts ballot $b_i = 1$, given policy platforms x_A and x_B. The probability of a vote for candidate B is then $1 - P_i(x_A, x_B)$. This representation of voting behavior allows for voters to vote in a deterministic fashion or, perhaps reflecting a lack of information on the candidates' parts, in a probabilistic way. Finally, let $P(x_A, x_B)$ denote the probability that candidate A wins the election, given the individual vote probabilities. Of course, B's probability of winning is one minus this amount. We assume that candidates are expected utility maximizers, so that candidate A seeks to maximize $E[U_A(x_A, x_B, b_1, \ldots, b_n)]$, where the expectation is taken over vectors of ballots $(b_1, \ldots, b_n)$ with respect to the distribution induced by the $P_i(x_A, x_B)$ probabilities, and likewise for candidate B. If we "integrate out" the vector of ballots, we may write expected utilities as functions $EU_A(x_A, x_B)$ and $EU_B(x_A, x_B)$ of the candidates' policy positions alone.

Given these expected utilities, we may subject the electoral model to an equilibrium analysis to illuminate the locational incentives of the candidates. We say a pair (x_A^*, x_B^*) of policy platforms is an *equilibrium* if neither candidate can gain by unilaterally deviating, i.e., for all policies y, we have

$$EU_A(y, x_B^*) \leq EU_A(x_A^*, x_B^*) \quad \text{and} \quad EU_B(x_A^*, y) \leq EU_B(x_A^*, x_B^*).$$

A *mixed strategy* for a candidate is a probability distribution on X, representing the probabilities with which the candidate adopts various policy platforms. Given mixed strategies for each candidate, candidate A's payoff, for example, can be calculated by integrating over x_A with respect to A's distribution and over x_B with respect to B's distribution. A *mixed strategy equilibrium* is a pair of probability distributions, reflecting the possibility that neither candidate may be able to precisely predict his or her opponent, such that neither candidate can gain by deviating unilaterally.

3 The Downsian Model

We now specify voting behavior, following Downs (1957) and others, by assuming that voters vote in an essentially deterministic fashion as a function of the candidates' platforms: each voter simply votes for the candidate who offers the best political platform, voting in a random way (by flipping a fair coin) only when indifferent.

Downsian model. We assume that each voter i has preferences over policies represented by a strictly concave, differentiable utility function u_i (so u_i satisfies the conditions imposed on candidate utilities under policy motivation). Under these assumptions, voter i has a unique *ideal policy,* denoted $\tilde{x}_i$, which yields a strictly higher utility than any other policy. A common special case is that of quadratic utility, in which a voter cares only about the distance of a policy from his or her ideal policy, e.g., $u_i(x) = -||x - \tilde{x}_i||^2$ for all x (where $||\cdot||$ is the usual Euclidean norm). In this case, the voters' indifference curves take the form of concentric circles centered at the ideal policy. For convenience, we assume that the ideal policies of the voters are distinct. We assume a voter votes for the candidate with the preferred policy position, randomizing only in case of indifference. This assumption is natural in this context and may be interpreted as "sincere" voting. Assuming voting is costless, it is also consistent with elimination of weakly dominated voting strategies in the voting game. Thus, we assume

$$P_i(x_A, x_B) \quad = \quad \begin{cases} 1 & \text{if } u_i(x_A) > u_i(x_B) \\ 0 & \text{if } u_i(x_A) < u_i(x_B) \\ \frac{1}{2} & \text{else.} \end{cases}$$

We will say that policy x is *majority-preferred* to policy y when more voters strictly prefer x than strictly prefer y. Writing $x \; M \; y$ to express this relation, we then formally have

$$x \; M \; y \quad \Leftrightarrow \quad \#\{i \mid u_i(x) > u_i(y)\} \; > \; \#\{i \mid u_i(y) > u_i(x)\}.$$

A closely related idea is that of the *core,* which is the set of maximal elements of this majority preference relation. That is, policy $\tilde{x}$ is in the core (or is a "core point") if there is no policy y such that $y \; M \; x$. Under our assumptions, the latter condition can be strengthened somewhat: if policy $\tilde{x}$ is in the core, then it is actually majority-preferred to every other policy y. Thus, there is at most one core point. Such a point is defined by the interesting normative and positive property that any move to a different policy will make some member of each majority coalition worse off.

Unfortunately, the core may be empty. The formal analysis of this issue relies on the concept of the gradient of a voter's utility function, the vector

$$\nabla u_i(x) \quad = \quad \left(\frac{\partial u_i}{\partial x_1}(x), \ldots, \frac{\partial u_i}{\partial x_d}(x) \right)$$

pointing in the direction of "steepest ascent" of the voter's utility at the policy x. Plott (1967) proved the following necessary and sufficient condition on voter gradients for a policy to belong to the core: any core point must be the ideal policy of some voter, and the gradients of the voters must be paired in such a way that, for every voter whose gradient points in one direction from the core point, there is a voter whose gradient points in the opposite direction. This condition, which is referred to as "radial symmetry," is satisfied in Figure 1, and it follows that voter 3's ideal point, $\tilde{x}_3$, in the figure is the core point.

Theorem 1 (Plott) *In the Downsian model, let policy $\tilde{x}$ be interior to the policy space X. If $\tilde{x}$ is the core point, then it is the ideal policy of some voter k and radial symmetry holds at $\tilde{x}$: each voter $i \neq k$ can be associated with a voter $j \neq k$ (in a 1-1 way) so that $\nabla u_j(\tilde{x})$ points in the direction opposite $\nabla u_i(\tilde{x})$.*

In accordance with this result, we refer to the voter k as the "core voter," and we denote the unique core point, when it exists, by $\tilde{x}_k$. The most common application of the Downsian model is the unidimensional model, where X is a subset of the real line and policies represent positions on a single salient issue. In this case, Theorem 1 implies that the only possible core point is the median of the voters' ideal policies, and since our assumptions imply that the preferences of all voters are single-peaked, Black's (1958) theorem implies that this median ideal policy is, indeed, the core point. Thus, in one dimension, the core is always non-empty and is characterized simply as the unique median ideal policy.

When the policy space is multidimensional, however, the necessary condition of radial symmetry becomes extremely restrictive — so restrictive that we would expect that the core is empty for almost all specifications of voter preferences.

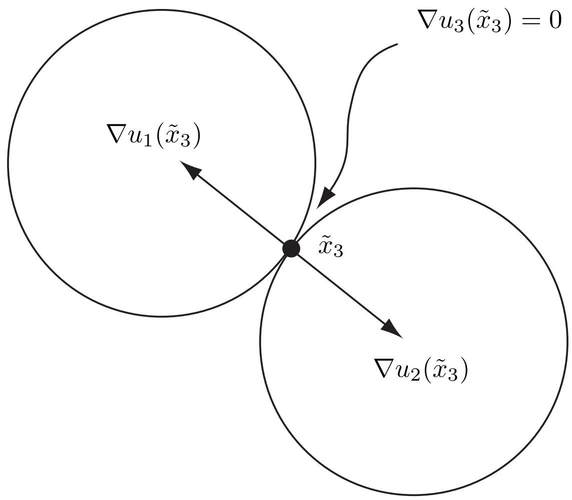

Figure 1: Radial symmetry

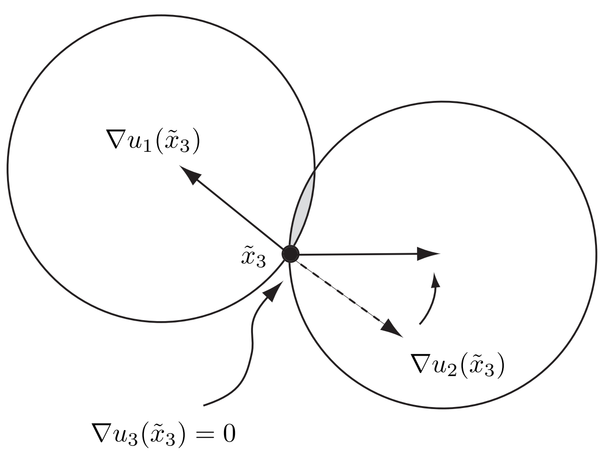

Figure 2: Violation of radial symmetry

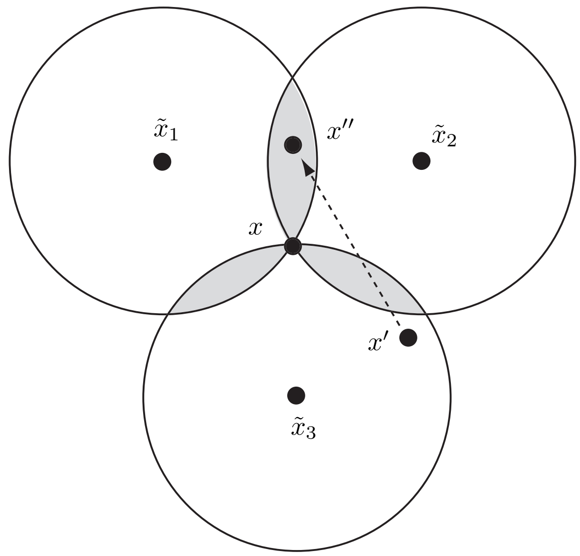

Figure 3: Discontinuity and non-convexity

Moreover, existence of a core point, if there is one, is a razor's edge phenomenon: if voter preferences were specified in such a way that the core was non-empty, then arbitrarily small perturbations of preferences could annihilate it. This is depicted in Figure 2, where the slightest move in voter 2's gradient breaks radial symmetry, making every policy in the shaded area majority-preferred to voter 3's ideal policy $\tilde{x}_3$. These ideas have been precisely formalized in the literature on the spatial model of social choice.[2]

3.1 Office Motivation

Under the assumption of deterministic voting, win motivation takes the following simple form:

$$EU_A(x_A, x_B) = \begin{cases} 1 & \text{if } x_A \; M \; x_B \\ 0 & \text{if } x_B \; M \; x_A \\ \frac{1}{2} & \text{else,} \end{cases}$$

and likewise for candidate B, while the alternative of vote motivation is

$$EU_A(x_A, x_B) = \#\{i \mid u_i(x_A) > u_i(x_B)\} + \frac{1}{2}\#\{i \mid u_i(x_A) = u_i(x_B)\},$$

and likewise for B.

[2]See, e.g., Cox (1984), Banks (1995), and Saari (1997).

It is easy to see that both objectives are marked by discontinuities and non-convexities, so that both of the conditions needed for the DFG theorem are violated. For example, Figure 3 depicts the indifference curves of three voters through a policy x, along with their ideal policies. If candidate B, say, locates at the policy x, then candidate A obtains a majority of votes by locating in any of the three shaded "leaves." The candidate's utility under win motivation from locating at x' is zero, and it is one at x''. Now consider A's utility when moving from x' directly to x'': when the candidate enters the shaded leaf, the utility jumps up discontinuously to one (violating continuity); it then jumps back down to zero and then back up to one (violating quasi-concavity) before reaching x''. Similar observations hold for vote motivation.

In the Downsian version of the electoral game, the distinction between the two formalizations of office motivation becomes irrelevant. Equilibria are completely characterized by the following result, which connects the concept of equilibrium from non-cooperative game theory to the notion of core from social choice theory, described above.

Theorem 2 *In the Downsian model, assume office motivation. There is an equilibrium* (x_A^*, x_B^*) *if and only if the core is non-empty. In this case, the equilibrium is unique, and the candidates locate at the core point:* $x_A^* = x_B^* = \tilde{x}_k$.

The argument for this result is elementary. As it is clear that both candidates locating at the core point is an equilibrium, I will prove only uniqueness of this equilibrium. Suppose (x_A^*, x_B^*) is an equilibrium, but one of the candidates, say B, locates at a policy not in the core. Then there is a policy z that beats x_B^* in a majority vote. Adopting z, candidate A can win the election with probability one and, of course, win more than $n/2$ votes. Since x_A^* is a best response to x_B^*, A's probability of winning at x_A^* must equal one under win motivation; likewise, A's expected vote must be greater than $n/2$ under plurality motivation. But then candidate B can deviate by locating at $x_B' = x_A^*$, winning with probability one half and obtaining an expected vote of $n/2$. Under either objective function, this deviation is profitable for B, a contradiction.

Theorem 2 has several important implications. First, the candidates must adopt identical policy positions in equilibrium. Second, this position is majority-preferred to all other policies, and so it is appealing on normative and positive grounds. Third, when the policy space is unidimensional, there is a unique equilibrium, and in equilibrium the candidates both locate at the median ideal policy. Known as the "median voter theorem," this connection was made by Hotelling (1929) in his model of spatial competition and by Downs (1957) in his classic analysis of elections.

Corollary 1 (Hotelling; Downs) *In the Downsian model, assume X is uni-*

dimensional and office-motivation. There is a unique equilibrium, and in equilibrium the candidates locate at the median ideal policy.

The pessimistic implication of Theorem 2, together with Theorem 1, is that equilibria of the multidimensional Downsian electoral game fail to exist for almost all specifications of voter preferences. A typical situation is depicted in Figure 3, where given an arbitrary location for the either candidate may profitably deviate to any policy in the three shaded leaves. And when voter preferences are such that equilibrium existence is obtained, it is susceptible to arbitrarily small perturbations of preferences. Returning to Figure 1, it is an equilibrium for both candidates to locate at $\tilde{x}_3$, where they each receive a payoff of one half. Perturbing voter 2's preferences, as in Figure 2, either candidate can deviate profitably to any policy in the shaded area to obtain a payoff of one.

An alternative is to look for equilibria in mixed strategies, which are modelled as probability distributions over the policy space and which allow for the possibility that one candidate may not be able to precisely predict the policy position of the other. Because the policy space is infinite in our model and the objective functions of the candidates are discontinuous, it is not known whether mixed strategy equilibria exist generally. One feature of the Downsian model that exacerbates these discontinuities is the inflexibility of voting behavior when a voter is indifferent between the positions of the candidates: the voter is assumed to vote for each candidate with equal probability. Duggan and Jackson (2005) show that, if we allow for indifferent voters to randomize with any probability between zero and one, then mixed strategy equilibria do indeed exist.[3]

Duggan and Jackson (2005) also show that in equilibrium, the support of the candidates' mixed strategies is contained in the "deep uncovered set," a centrally located subset of the policy space related to McKelvey's (1986) uncovered set. An implication of their results is that if voter preferences are specified so that a core point exists, and if we perturb voter preferences slightly, then the equilibrium mixed strategies of the candidates will put probability arbitrarily close to one on policies near the original core point. Thus, while the existence of pure strategy equilibria is knife-edge, mixed strategy equilibrium outcomes change in a continuous way when voter preferences are perturbed.

[3]Duggan and Jackson (2005) do require that indifferent voters treat the candidates symmetrically. If voter i is indifferent between x_A and x_B and votes for candidate A with probability, say α, then i must vote for A with probability $1 - \alpha$ if the candidates switch positions.

3.2 Policy Motivation

The objective function representing policy motivation takes the following form under deterministic voting:

$$EU_A(x_A, x_B) \quad = \quad \begin{cases} u_A(x_A) & \text{if } x_A \; M \; x_B \\ u_A(x_B) & \text{if } x_B \; M \; x_A \\ \frac{u_A(x_A)+u_A(x_B)}{2} & \text{else.} \end{cases}$$

The above objective function also suffers from discontinuities and non-convexities, but it differs fundamentally from the case of office motivation in other ways. Nevertheless, Roemer (1994) proves a median voter theorem for policy motivation, and Wittman (1977) and Calvert (1985), prove related results for multidimensional policy spaces and Euclidean preferences that yield the median voter result in the special case of one dimension. As with office motivation, there is a unique equilibrium, and in equilibrium both candidates locate at the median ideal policy.

Theorem 3 (Wittman; Calvert; Roemer) *In the Downsian model, assume X is unidimensional and policy motivation. If $\tilde{x}_A < \tilde{x}_k < \tilde{x}_B$, then there is a unique equilibrium (x_A^*, x_B^*). In equilibrium, the candidates locate at the median ideal policy: $x_A^* = x_B^* = \tilde{x}_k$.*

Theorem 3 leaves open the question of whether equilibria exist when the policy space is multidimensional. There is some reason to believe that the negative result of Theorem 2 might be attenuated in the case of pure policy motivation: under win motivation, a candidate could benefit from a move to *any* policy position that is majority-preferred to that of his or her opponent, creating a large number of potential profitable deviations; under policy motivation, however, a candidate can only benefit from moving to a policy position that he or she prefers to her opponent's. This restricts the possibilities for profitable deviations and improves the prospects for finding equilibria where none existed under win motivation. This possibility is noted by Duggan and Fey (2005a) and is depicted in Figure 4. In this example, the core is empty and there is no equilibrium under win motivation, yet it is an equilibrium under policy motivation for the candidates to locate at voter 3's ideal policy. As the candidates' indifference curves suggest, neither can move to a position that is both preferable to $\tilde{x}_3$ and beats $\tilde{x}_3$ in a majority vote. Moreover, existence of equilibrium in this example is robust to small perturbations in the preferences of voters and candidates.

In Figure 4, we specify that candidates A and B choose identical policy positions. While this form of policy convergence is well-known from the unidimensional model, Duggan and Fey (2005a) show that it is a near universal feature of electoral competition with policy motivated candidates, regardless of

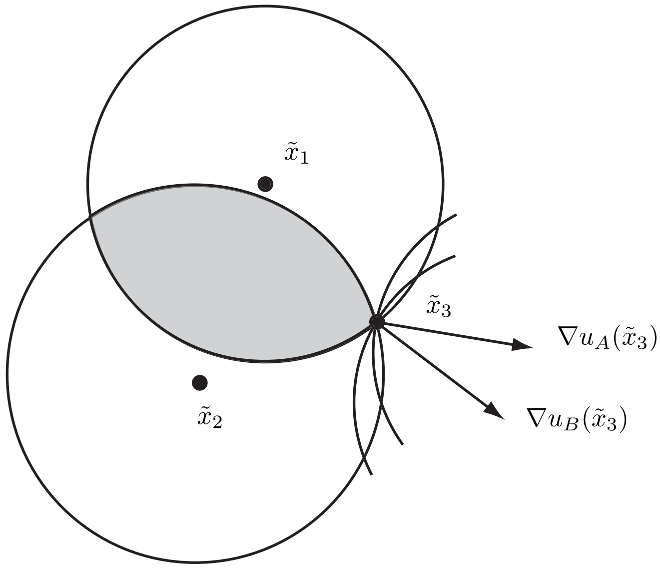

Figure 4: Equilibrium with policy motivation but not win motivation

the dimensionality of the policy space. Roemer (2001) also argues that equilibria in which the two candidates adopt distinct policy positions almost never exist.

Theorem 4 (Duggan and Fey) *In the Downsian model, assume policy motivation. If (x_A^*, x_B^*) is an equilibrium such that neither candidate locates at his or her ideal policy, i.e., $\nabla u_A(x_A^*) \neq 0$ and $\nabla u_B(x_B^*) \neq 0$, then the candidates' policy positions are identical: $x_A^* = x_B^* = x^*$.*

The next result shows that the kind of positive result depicted in Figure 4 is limited to the two-dimensional case. The result, due to Duggan and Fey (2005a), gives a strong necessary condition on equilibria at which the candidates' gradients do not point in the same direction, as they do not in Figure 4. Like radial symmetry from Plott's theorem, this necessary condition requires that the gradients of certain voters be diametrically opposed, though now the restriction applies only to voters whose gradients do not lie on the plane spanned by the gradients of the candidates.

Theorem 5 (Duggan and Fey) *In the Downsian model, assume policy motivation. If (x_A^*, x_B^*) is an equilibrium of the electoral game such that $x_A^* = x_B^* = x^*$, and if the candidates' gradients at x^* do not point in the same direction, then x^* is the ideal policy of some voter k, i.e., $x^* = \tilde{x}_k$; and each voter $i \neq k$ whose gradient does not lie on the plane spanned by the candidates' gradients* *can be associated with a voter $j \neq k$ (in a 1-1 way) so that $\nabla u_j(x^*)$ points in the direction opposite $\nabla u_i(x^*)$.*

The implications of this theorem are sharpest when the policy space has dimension at least three: then the plane spanned by the candidates' gradients is lower-dimensional, and there will typically be at least one voter whose gradient does not lie on this plane; and this voter must be exactly opposed by another. We conclude that equilibria will almost never exist, and if there is an equilibrium, then existence will necessarily be susceptible to even small perturbations of voter or candidate preferences. Thus, while policy motivation restricts the set of potential profitable moves, multidimensional policy spaces offer the candidates sufficient scope for deviations that equilibria will typically fail to exist. As in the case of office motivation, however, the result of Duggan and Jackson (2005) yields existence of a mixed strategy equilibrium when we allow indifferent voters to randomize in a more flexible way.

4 Probabilistic Voting: The Stochastic Partisanship Model

In the Downsian model, we assume that voters behave in a deterministic fashion (unless indifferent between the candidates) and that the candidates can predict voting behavior precisely. The literature on probabilistic voting relaxes these assumptions, viewing the ballots of voters as random variables. While this class of models may capture indeterminacy inherent in the behavior of voters, it is also consistent with the rational choice approach: it may be that the decision of a voter as ultimately determined by the voter's preferences, but we allow for the possibility that the candidates do not perfectly observe the preferences of voters; instead, candidates have beliefs about the preferences of voters, and therefore their behavior, and we model these beliefs probabilistically.

In contrast to the Downsian model, it now becomes important to distinguish between the two types of office motivation, as reflected in the results surveyed below. In addition, we must be explicit about the structure of the voters' decision problems and the source of randomness in the model. The approach we consider in this section endows voters with policy preferences that are known to the candidates (and therefore taken as given), but it assumes that the voters also have partisan preferences over the candidates unrelated to their policy positions. The intensities of these partisan preferences are unknown to the candidates.

Stochastic partisanship model. Assume that each voter i has a strictly concave, differentiable utility function u_i, as in the deterministic voting model, but now assume a "utility bias" β_i in favor of candidate B.[4] We incorporate these biases into our model of voting behavior by assuming that i votes for A if and only if the utility of candidate A's platform exceeds that of B's platform by at least β_i. That is, i votes for A if and only if $u_i(x_A) \geq u_i(x_B) + \beta_i$, i.e., $\beta_i \leq u_i(x_A) - u_i(x_B)$. We assume that the profile $(\beta_1, \ldots, \beta_n)$ of biases is a random variable from the candidates' perspective, and we assume that each β_i is distributed (not necessarily independently) according to the distribution F_i. We assume that each F_i is continuous and strictly increasing on an interval that includes all possible utility differences $u_i(x) - u_i(y)$, as x and y range over all of X,[5] so that the probability that voter i votes for candidate A given platforms x_A and x_B is $F_i(u_i(x_A) - u_i(x_B))$. We do not assume biases are identically distributed, but we impose a weak symmetry assumption. Letting f_i denote the density of F_i, we assume for simplicity that the likelihood that a voter is unbiased is the same for all voters: $f_i(0) = f_j(0)$ for all voters i and j. A convenient special case is that in which the F_i are uniform with identical supports. We refer to this as the *uniform partisanship model.*

We first consider the known results on equilibrium existence and characterization under the assumption of vote-motivated candidates, and we then examine the results for win motivation.

4.1 Vote Motivation

In the general probabilistic voting framework, vote motivation for the candidates can be written

$$EU_A(x_A, x_B) \quad = \quad \sum_i P_i(x_A, x_B),$$

with $EU_B(x_A, x_B)$ equal to n minus the above quantity. Because F_i and u_i are assumed continuous, it follows that candidate A's utility, $\sum_i F_i(u_i(x_A) - u_i(x_B))$ is continuous, and likewise for B. Furthermore, as we will see, the linear form of the candidates utilities in terms of individual vote probabilities invites a simple sufficient condition under which quasi-concavity, the second condition of the DFG theorem, holds.

In the stochastic partisanship model, denote the unique maximizer of the sum of voter utilities by

$$\overline{x} \quad = \quad \arg\max_{x \in X} \sum_i u_i(x).$$

[4] In case β_i is negative, we can think of this as a bias for candidate A. We maintain the terminology with candidate B as the point of reference.

[5] An implication is that, given any two platforms for the candidates, there is some chance (perhaps very small) that the voter's bias outweighs any policy considerations.

This policy is often referred to as the "utilitarian optimum," though this term suggests welfare connotations that are difficult to justify.[6] We use the somewhat more neutral term *utilitarian point.* When voter utilities are quadratic, it is well-known that the utilitarian point is equal to the mean of the voters' ideal policies.

Our first result establishes that, in equilibrium, the candidates must offer voters the same policy position. This policy is exactly the utilitarian point, implying that the candidates must adopt the same central position in the policy space in equilibrium, regardless of the dimensionality of the policy space. Versions of this result have appeared in several places, notably in the work of Hinich (1977, 1978) and Lindbeck and Weibull (1987, 1993).[7] The general statement here is due to Banks and Duggan (2005).[8] When voter utilities are quadratic, an implication is Hinich's "mean voter" theorem.

Theorem 6 (Hinich; Lindbeck and Weibull; Banks and Duggan) *In the stochastic partisanship model, assume vote motivation. If (x_A^*, x_B^*) is an interior equilibrium, then both candidates locate at the utilitarian point: $x_A^* = x_B^* = \overline{x}$.*

To prove the result, consider any interior equilibrium (x_A^*, x_B^*). By Theorem 8 of Banks and Duggan (2005), the candidates must adopt identical platforms, i.e., $x_A^* = x_B^* = x^*$ for some policy x^*. Candidate A's maximization problem is then

$$\max_{x \in X} \sum_i F_i(u_i(x) - u_i(x^*)),$$

and the first order condition at the equilibrium platform $x = x^*$ is

$$\sum_i f_i(0) \nabla u_i(x^*) \quad = \quad 0.$$

Since we assume equal likelihoods of zero bias for the voters, this reduces to $\nabla \sum_i u_i(x^*) = 0$. Since $\sum_i u_i$, as the sum of strictly concave functions, is strictly concave, it follows that x^* maximizes $\sum_i u_i$, i.e., $x^* = \overline{x}$.

The next result states known conditions for existence of an equilibrium. Since the objective functions of the candidates are continuous, the key is to ensure that quasi-concavity is satisfied. The sufficient conditions we give are

[6] Note that an individual's vote probability P_i and a distribution F_i pin down a unique utility function u_i in the stochastic partisanship model. Our symmetry requirement that $f_i(0) = f_j(0)$ then allows us to compare voter utilities, but there is no special normative basis for this.

[7] See also Coughlin (1992). Ledyard (1984) derives a similar utilitarian result from a model of costly and strategic voting.

[8] Hinich (1978) assumes that, given the same utility difference for the candidates, any two voters will have the same marginal vote propensities. Lindbeck and Weibull (1993) assume strict quasi-concavity of candidate payoffs and, implicitly, symmetry of the electoral game.

fulfilled, for example, if the bias terms of the voters are uniformly distributed, or if the distributions F_i are "close enough" to uniform. Thus, in contrast to Theorem 2, which implies the generic non-existence of equilibria in multiple dimensions under deterministic voting, Theorem 7 offers reasonable (if somewhat restrictive) conditions that guarantee an equilibrium under probabilistic voting. Hinich, Ledyard, and Ordeshook (1972, 1973) give similar sufficient conditions in a model that allows for abstention by voters, and Enelow and Hinich (1989) and Lindbeck and Weibull (1993) make similar observations.[9] By Theorem 6, if there is an equilibrium, then it is unique, and both candidates locate at the utilitarian point.

Theorem 7 (Hinich, Ledyard, and Ordeshook; Lindbeck and Weibull) *In the stochastic partisanship model, assume vote motivation, and assume the following for each voter i:*

- $F_i(u_i(x))$ *is concave in* x
- $F_i(-u_i(x))$ *is convex in* x.

There exists an equilibrium.

The proof of Theorem 7 is straightforward. Continuity of the candidates' expected utilities has already been noted. By the assumptions of the proposition, $P_i(x_A, x_B) = F_i(u_i(x_A) - u_i(x_B))$ is a concave function of x_A. Therefore, as the sum of concave functions, $EU_A(x_A, x_B)$ is a concave function of x_A, and a similar argument holds for candidate B. Thus, existence of an equilibrium follows from the DFG theorem.

Theorems 6 and 7 taken together may suggest a puzzling discrepancy between the Downsian and stochastic partisanship models. Consider the possibility of modifying the Downsian model by introducing a "small" amount of bias, i.e., consider distributions F_i of biases that converge to the point mass on zero.[10] In this way, we can satisfy the assumptions of the stochastic partisanship model in models arbitrarily close to the Downsian model. By Theorem 6, the equilibria of these stochastic partisanship models must be at the utilitarian point, and it may therefore appear that the equilibrium moves from the median ideal policy in the Downsian model to the utilitarian point in the presence of the slightest noise in voting behavior. Or, to use Hinich's (1977) terminology, it may appear that the median is an "artifact."

[9] Note that Lindbeck and Weibull (1993) implicitly rely on the assumption that the electoral game is symmetric in their Theorem 1, an assumption not made here.

[10] Here, "convergence" is in the sense of weak* convergence, a convention we maintain when referring to probability distributions.

In fact, however, Theorem 6 only gives a necessary condition for equilibria in the stochastic partisanship model: it says that *if* there is an equilibrium, then both candidates must locate at the utilitarian point, leaving the possibility that there is no equilibrium. Banks and Duggan (2005) and Laussel and Le Breton (2002) show that this is necessarily the case: when voting behavior is close to deterministic in the stochastic partisanship model, there is *no* equilibrium in pure strategies.[11] Thus, the introduction of probabilistic voting into the Downsian model can actually create equilibrium existence problems, even in the unidimensional model, where the median voter theorem holds.

In contrast, since only continuity of the candidates' utilities is required for existence of a mixed strategy equilibrium, we have general existence of mixed strategy equilibria in the stochastic partisanship model under no additional assumptions.

Theorem 8 *In the stochastic partisanship model, assume vote motivation. There is a mixed strategy equilibrium.*

Although pure strategy equilibria will not exist when we add a small amount of noise to voting behavior in the Downsian model, Theorem 8 implies that there will be mixed strategy equilibria. Moreover, Banks and Duggan (2005) prove that these mixed strategy equilibria must converge to the median ideal policy as the amount of noise goes to zero, and we conclude that when a small amount of noise is added to voting behavior in the Downsian model, the equilibrium does not suddenly move to the utilitarian point. Pure strategy equilibria cease to exist, but mixed strategy equilibria do exist, and policies close to the median will be played with probability arbitrarily close to one in these equilibria.

Our results for the stochastic partisanship model have immediate consequences for a closely related model in which voter biases enter into voting behavior in a multiplicative way.

Stochastic multiplicative bias model. Assume each voter i has a positive, strictly log-concave, smooth utility function u_i and a positive utility bias β_i in favor of candidate B. In contrast to the stochastic partisanship model, assume that i votes for A if and only if $u_i(x_A) \geq u_i(x_B)\beta_i$. As before, each β_i is distributed according to a distribution F_i. We assume each F_i is continuous and strictly increasing on an interval that includes all ratios of utilities $u_i(x)/u_i(y)$, so that the probability voter i votes for candidate A is $F_i(u_i(x_A)/u_i(x_B))$. Following the partisanship model, assume $f_i(1) = f_j(1)$ for all voters i and j. Let

$$\hat{x} \quad = \quad \arg\max_{x \in X} \prod_i u_i(x)$$

[11] This overturns Theorem 2 of Hinich (1978). See Banks and Duggan (2005) for an extended discussion.

maximize the product of voter utilities. This welfare function is known as the "Nash welfare" function, and we refer to $\hat{x}$ as the *Nash point.*

The next result is proved by Coughlin and Nitzan (1982) for the special case of the binary Luce model. The general statement here is due to Banks and Duggan (2005).

Corollary 2 (Coughlin and Nitzan; Banks and Duggan) *In the stochastic multiplicative bias model, assume vote motivation. If (x_A^*, x_B^*) is an interior equilibrium, then both candidates locate at the Nash point: $x_A^* = x_B^* = \hat{x}$.*

The proof follows directly from Theorem 6, after an appropriate transformation of the stochastic multiplicative bias model. Given a model satisfying the assumptions of the corollary, define an associated partisanship model as follows:

$$\begin{aligned}\dot{u}_i(x) &= \ln(u_i(x)) \\ \dot{F}_i(u) &= F_i(e^u).\end{aligned}$$

Note that the individual vote probabilities in the two models are identical, as

$$\begin{aligned}\dot{P}_i(x_A, x_B) \;&=\; \dot{F}_i(\dot{u}_i(x_A) - \dot{u}_i(x_B)) \;=\; \dot{F}_i(\ln(u_i(x_A)/u_i(x_B))) \\ &=\; F_i(u_i(x_A)/u_i(x_B)) \;=\; P_i(x_A, x_B).\end{aligned}$$

and so, therefore, the candidates' objective functions are identical as well. If (x_A^*, x_B^*) is an equilibrium in the stochastic multiplicative bias model, it follows that it is an equilibrium in the associated partisanship model. By Theorem 6, applied to the associated partisanship model, we therefore have

$$x_A^* \;=\; x_B^* \;=\; \arg\max_{x \in X} \sum_i \dot{u}_i(x) \;=\; \arg\max_{x \in X} \prod_i u_i(x) \;=\; \hat{x},$$

as required.

By a similar logic, we can extend Theorem 7 to the stochastic multiplicative bias model as well, deriving sufficient conditions for existence of an equilibrium of the corresponding electoral game: it is enough if the distributions F_i, composed with the exponential function, possess the appropriate convexity properties. Note the expositional tradeoff apparent in Corollaries 2 and 3: we use only log-concavity of voter utility functions for the characterization of equilibrium at the Nash point, weakening the assumption of concavity in Theorem 6, but the conditions on distributions needed for equilibrium existence in the stochastic multiplicative bias model are more restrictive than those of Theorem 7.

Corollary 3 (Coughlin and Nitzan) *In the stochastic multiplicative bias model, assume vote motivation, and assume the following for each voter i:*

- $F_i(e^{u_i(x)})$ *is concave*
- $F_i(e^{-u_i(x)})$ *is convex.*

There exists an equilibrium.

For an example of a distribution fulfilling the requirements of Corollary 3, consider the binary Luce model used by Coughlin and Nitzan (1981), where

$$P_i(x_A, x_B) \quad = \quad \frac{u_i(x_A)}{u_i(x_A) + u_i(x_B)}.$$

This is the special case of the stochastic multiplicative bias model with $F_i(u) = \frac{u}{u+1}$, and it is then apparent that

$$F_i(e^{u_i(x)}) \quad = \quad \frac{1}{1 + e^{-u_i(x)}}$$

is concave in x. These arguments illustrate that the stochastic partisanship and multiplicative bias models are equivalent, up to a simple transformation. We could just as well have begun with Corollaries 2 and 3 and derived Theorems 6 and 7 for the stochastic partisanship model.

4.2 Win Motivation

In the general probabilistic voting framework, win motivation for the candidates takes the form

$$EU_A(x_A, x_B) \quad = \quad P(x_A, x_B),$$

with candidate B's utility equal to one minus the above quantity. Because the candidates' probability of winning is the probability of a particular event (receiving more votes than the opponent) with respect to a binomial probability distribution, this objective function lacks the nice linear form of vote motivation.

As under vote motivation, the objective functions of the candidates are continuous. We will see that quasi-concavity is more difficult to maintain under win motivation, but there is, nevertheless, a close connection between the equilibria generated by the two objective functions. Our first result, due to Duggan (2000a), establishes that there is only one possible equilibrium under win motivation: the utilitarian point, familiar from the analysis of vote motivation. Thus, again, equilibrium incentives drive the candidates to take identical positions at a central point in the policy space.

Theorem 9 (Duggan) *In the stochastic partisanship model, assume win motivation, and assume the following for each voter i:*

- $F_i(0) = 1/2$
- *$F_i(u_i(x))$ is concave in x*
- *$F_i(-u_i(x))$ is convex in x.*

If (x_A^*, x_B^*) is an interior equilibrium, then both candidates locate at the utilitarian point: $x_A^* = x_B^* = \overline{x}$.

Of the conditions imposed on individual vote probabilities, the first merely assumes a weak form of symmetry, so that neither candidate is ex ante advantaged. The second and third conditions are those used in Theorem 7 to ensure the existence of an equilibrium under vote motivation. Here, those conditions serve a related purpose, which is depicted graphically in Figure 5 for the case of two voters. Fixing candidate B's policy at the utilitarian point, $x_B = \overline{x}$, a choice of x_A will determine a vector

$$(P_1(x_A, \overline{x}), P_2(x_A, \overline{x}), \ldots, P_n(x_A, \overline{x}))$$

of individual vote probabilities for candidate A. Varying x_A within the policy space X, this generates the set of probability vectors "achievable" by candidate A, the shaded area in Figure 5. Note that, by the symmetry assumption of Theorem 9, candidate A wins with probability one half if the candidate chooses $x_A = \overline{x}$, and so the vector $(1/2, 1/2)$ lies in this set. The second and third assumptions in Theorem 9 imply that this set of achievable probability vectors is convex.

A consequence of Theorem 9 is that, in order to obtain a full understanding of equilibria under win motivation, we need only understand the conditions under which it is indeed an equilibrium for the candidates to locate at the utilitarian point. Clearly, under the conditions of Theorem 7, locating at the utilitarian point is an equilibrium under vote motivation: this is apparent in Figure 5, as the tangency at probability vector $(1/2, 1/2)$ shows that there is no other achievable vector of probabilities lying on a higher level set for expected vote. In contrast, the probability-of-winning objective is not linear in vote probabilities, and its convexity properties are generally poor. This is illustrated figuratively in Figure 5 by the "dip" in the probability-of-winning level set, which creates the possibility for a deviation by candidate A that is profitable under win motivation.

The next result gives a simple condition in terms of the second derivatives of voters' utility functions that rules out the problem illustrated in Figure 5. Assuming, for simplicity, that individual bias terms are uniformly distributed, it is sufficient to assume that the Hessian matrix of at least one voter's utility function is negative definite at the utilitarian point. Since we maintain the assumption that utilities are concave, the Hessian is already negative semi-definite,

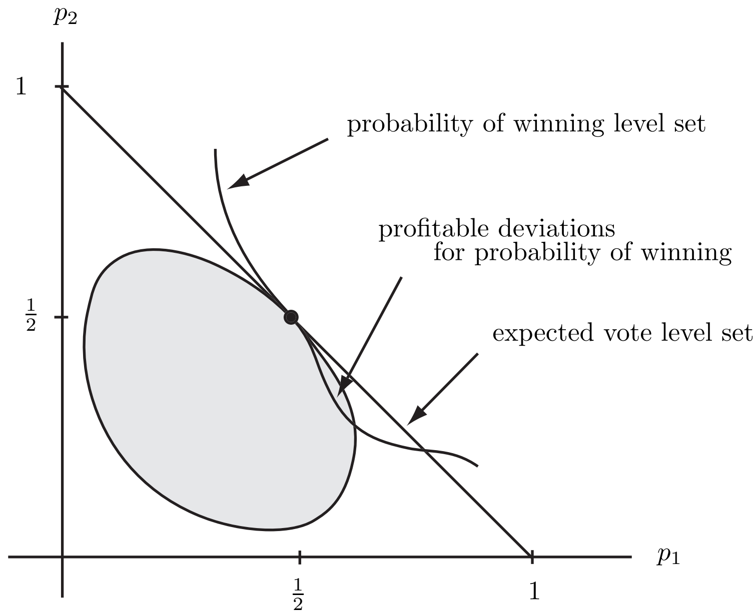

Figure 5: Difficulty with win motivation

so the added restriction of negative definiteness seems hardly objectionable. The result does not deliver the existence of a global equilibrium, but only one that is immune to small deviations by one of the candidates. The statement here is in the context of the uniform partisanship model used in Duggan (2000a), whereas Patty (2003) provides an extension to a more general model of voting and to multiple candidates.[12] Note that the result does not deliver the existence of a (global) equilibrium, but only a "local" equilibrium, which is immune to small deviations by one of the candidates. Thus, it leaves the possibility that one candidate could increase his or her probability of winning by positioning far from the utilitarian point.

Theorem 10 (Duggan; Patty) *In the uniform partisanship model, assume win motivation. If the Hessian matrix of u_i at $\overline{x}$ is negative definite for some voter i, then $(\overline{x}, \overline{x})$ is a local equilibrium.*

Interestingly, the added restriction of negative definiteness is needed for Theorem 10: the assumption of strict concavity alone is not enough for the result. Duggan (2000a) gives the following three-voter, unidimensional example of the uniform partisanship model in which the utilitarian point is not a local equilib-

[12]See Aranson, Hinich, and Ordeshook (1974) for earlier results on the equivalence of vote and win motivation.

rium:

$$\begin{aligned} u_1(x) &= -2x - 4(x-1/2)^4 \\ u_2(x) &= x - 2(x-1/2)^4 \\ u_3(x) &= u_2(x). \end{aligned}$$

It is straightforward to verify that 1's, 2's, and 3's ideal points are 0, 1, and 1, respectively, and that $\overline{x} = 1/2$. Thus, $(1/2, 1/2)$ is the unique equilibrium under vote motivation. These utility functions are all strictly concave, but we have $u_1''(\overline{x}) = u_2''(\overline{x}) = u_3''(\overline{x}) = 0$, violating the negative definiteness requirement in Theorem 10. In fact, arbitrarily small increases in A's platform will produce a probability of winning greater than one half, and we conclude that $(1/2, 1/2)$ is not a local equilibrium under probability of winning.

In contrast, since the candidates' objective functions are continuous in the stochastic partisanship model, the DFG theorem yields the general existence of a mixed strategy equilibrium.

Theorem 11 *In the stochastic partisanship model, assume win motivation. There is a mixed strategy equilibrium.*

Furthermore, results of Kramer (1978) and Duggan and Jackson (2005) show that the mixed strategy equilibria in models close to the Downsian model must put probability arbitrarily close to one on policies close to the median ideal policy. As with vote motivation, the equilibrium policy locations of the candidates change in a continuous way when noise is added to the Downsian model.

5 Probabilistic Voting: The Stochastic Preference Model

The second main approach to modelling probabilistic voting focuses only on policy considerations (dropping considerations of partisanship) and allows for the possibility that the candidates do not perfectly observe the policy preferences of voters.

Stochastic preference model. Assume each voter i has a strictly concave, differentiable utility function $u_i(\cdot, \theta_i)$, where θ_i is a preference parameter lying in a Euclidean space Θ. We assume that the vector $(\theta_1, \ldots, \theta_n)$ of parameters is a random variable from the candidates' perspective, and we assume that each θ_i is distributed according to a distribution function G_i. We do not assume that these random variables are independent, but we assume that the distribution of preferences is sufficiently "dispersed" for each voter, in the following sense: for every voter i and all distinct policies x and y, we have

$$\Pr(\{\theta_i \mid u_i(z,\theta_i) > u_i(y,\theta_i) > u_i(x,\theta_i)\}) \quad > \quad 0,$$

where z is the midpoint between x and y, and we also have

$$\Pr(\{\theta_i \mid u_i(x,\theta_i) = u_i(y,\theta_i)\}) \quad = \quad 0.$$

Given parameter θ_i, our assumptions imply that voter i has a unique ideal policy, and we let H_i denote the distribution of voter i's ideal policy. In the case of a unidimensional policy space, for example, $H_i(x)$ is the probability that i's ideal policy is less than or equal to x, and by our dispersion assumption it is continuous and strictly increasing. It is common to identify θ_i with voter i's ideal policy and to assume that $u_i(\cdot,\theta_i)$ is quadratic, and our dispersion condition is then satisfied under the uncontroversial assumption that the distribution of ideal policies is continuous with full support.

When all voters appear *ex ante* identical to the candidates, we drop the i index on u and G, and we assume without loss of generality that there is a single voter. We refer to this as the *representative voter stochastic preference model.* In the quadratic model with a unidimensional policy space, we may focus on the representative voter model without loss of generality. To see this, given a realization $(\theta_1,\ldots,\theta_n)$ of ideal policies, let θ_k denote the median ideal policy. It is well-known that the median voter k is "decisive," in the sense that a candidate wins a majority of the vote if and only if his or her policy position is preferred by voter k to the other candidate's position. Thus, we have

$$P(x_A,x_B|\theta_1,\ldots,\theta_n) \quad = \quad \begin{cases} 1 & \text{if } |x_A-\theta_k| < |x_B-\theta_k| \\ 0 & \text{if } |x_B-\theta_k| < |x_A-\theta_k| \\ \frac{1}{2} & \text{else.} \end{cases}$$

We can capture this formally by assuming a single voter and letting G be the distribution of the median ideal policy in the n-voter model, a special case we call the *quadratic stochastic preference model.*

Paralleling the previous section, we first consider the objective of vote motivation, and we then consider win motivation, both objectives defined as before. We end with the analysis of policy motivation, which has been considered primarily in the context of the stochastic preference model.

5.1 Vote Motivation

An immediate technical difference between the stochastic preference and stochastic partisanship models is that we now lose full continuity of the candidates' expected votes, as discontinuities appear along the "diagonal," where $x_A = x_B$.

To see this, consider the quadratic stochastic preference model in the context of a unidimensional policy space. Fix x_B to the right of the median of G, and let x_A approach x_B from the left. Then candidate A's expected utility converges to $G(x_B) > 1/2$, but at $x_A = x_B$, A's expected utility is $1/2$. The only value of x_B where such a discontinuity does not occur is at the median of G. Thus, the stochastic preference model generally exhibits discontinuities when one candidate "crosses over" the other.

Despite the presence of these discontinuities, there is a unique equilibrium when the policy space is unidimensional, and it is easily characterized. Let H_α be the distribution defined by $H_\alpha(x) = \frac{1}{n}\sum_i H_i(x)$, and let x_α be the unique median of this average distribution. The next result, which is proved in Duggan (2005), establishes that in equilibrium the candidates must locate at the same policy, the median x_α of the average distribution. Thus, we are back to a "median-like" result, but now the equilibrium is at the median of the average distribution.

Theorem 12 (Duggan) *In the stochastic preference model, assume X is unidimensional and vote motivation. There is a unique equilibrium (x_A^*, x_B^*). In equilibrium, both candidates locate at the median of the average distribution: $x_A^* = x_B^* = x_\alpha$.*

It is relatively easy to prove that it is an equilibrium for both candidates to locate at x_α, so I will prove only uniqueness. Suppose that in equilibrium some candidate locates at a policy x^* other than x_α. Without loss of generality, assume $x^* < x_\alpha$, so that $H_\alpha(x) < 1/2$. Since the electoral game is symmetric and zero-sum, a standard interchangeability argument shows that it must be an equilibrium for both candidates to locate at x^*, where the expected utility of each candidate is $n/2$. But allow candidate A to deviate by moving slightly to the right of x^* to a position $x^* + \epsilon < x_\alpha$. Since a voter with an ideal policy to the right of $x^* + \epsilon$ will vote for A, we have

$$\begin{aligned} EU_A(x^* + \epsilon, x^*) &\geq \sum_i \left[1 - H_i\left(x^* + \epsilon\right)\right] \\ &= n - \sum_i H_i\left(x^* + \epsilon\right). \end{aligned}$$

By the assumption that $x^* + \epsilon < x_\alpha$, this quantity is greater than $n/2$. Therefore, candidate A can achieve an expected vote greater than $n/2$ by deviating from x^*, a contradiction. This argument shows that the only possible equilibrium is for both candidates to locate at x_α.

Now consider a stochastic preference model close to the Downsian model, in the sense that each distribution G_i piles probability mass near some value of θ_i. Assuming a unidimensional policy space, the median voter theorem implies that, in the Downsian model where voters' preferences are given by $u_i(\cdot, \theta_i)$, the unique equilibrium is for both candidates to locate at the median ideal policy. By Theorem 12, the stochastic preference model admits a unique equilibrium, and in equilibrium the candidates locate at the median of the average distribution. As the model gets close to Downsian, the average distributions will approach the distribution of ideal policies in the Downsian model, and the equilibrium points x_α will approach the median ideal policy. Therefore, when we add a small amount of noise to voting behavior in the Downsian model, the equilibrium changes in a continuous way.

When the policy space is multidimensional, the task of equilibrium characterization is more difficult, and we simplify matters by specializing to the representative voter stochastic preference model. We say a policy x is a *generalized median in all directions* if, compared to every other policy y, the voter is more likely to prefer x to y than the converse: for every policy y, we have

$$\Pr(\{\theta \mid u(y,\theta) > u(x,\theta)\}) \quad \leq \quad \frac{1}{2}.$$

By strict concavity and our dispersion condition, if there is a generalized median in all directions, then there is exactly one, which we denote x_γ.[13] In the quadratic version of the model, x_γ is equivalent to a median in all directions, in the usual sense.[14] When the policy space is multidimensional, such a policy exists, for example, if G has a radially symmetric density function, such as the normal distribution. This can be weakened, but existence of a median in all directions is generally quite restrictive when the policy space has dimension at least two.

The next result provides a characterization of equilibria in the stochastic preference model: in equilibrium, the candidates must locate at the generalized median in all directions. Under our maintained assumptions, a generalized median in all directions is essentially an "estimated median," so the next result is very close to a result due to Calvert (1985), and it is similar in spirit to results of Davis and Hinich (1968) and Hoyer and Mayer (1974). As we have seen before, strategic incentives drive the candidates to take identical positions in equilibrium, but the implication for equilibrium existence in multidimensional policy spaces is negative, as existence of a generalized median in all directions is extremely restrictive.

[13]Suppose there were distinct generalized medians in all directions, x and y, with midpoint z. By strict concavity, if $u(x,\theta) > u(y,\theta)$, then $u(z,\theta) > u(y,\theta)$. Therefore, since x is a generalized median in all directions, the probability the voter strictly prefers z to y is at least one half. By dispersion, there is also positive probability that $u(z,\theta) > u(y,\theta) > u(x,\theta)$. Therefore, the probability the voter strictly prefers z to y is greater than one half, contradicting the assumption that y is a generalized median in all directions.

[14]In two dimensions, this means that every line through x_γ divides the space in half: the probability that the representative voter's ideal policy is to one side of the line is equal to one half.

Theorem 13 (Calvert) *In the representative voter stochastic preference model, assume vote motivation. There is an equilibrium* (x_A^*, x_B^*) *if and only if there is a generalized median in all directions. In this case, the equilibrium is unique, and the candidates locate at the generalized median point:* $x_A^* = x_B^* = x_\gamma$.

To prove the theorem, suppose (x_A^*, x_B^*) is an equilibrium, but one of the candidates, say B, does not locate at the generalized median policy. Then there is some policy z such that each voter prefers z to x_B with probability strictly greater than one half. Thus, adopting z, candidate A can win the election with probability greater than one half, and since x_A^* is a best response to x_B^*, A's probability of winning at x_A^* can be no less. But then candidate B can deviate by locating at $x'_B = x_A^*$, matching A and winning with probability one half. Since this increases B's probability of winning, the initial pair of policy positions cannot be an equilibrium, a contradiction.

Because discontinuities in the stochastic preference model are restricted to the diagonal, where the candidates choose identical platforms, existence of mixed strategy equilibria is more easily obtained than in the Downsian model. In this case, the results of Duggan and Jackson (2005) yield an equilibrium with no modifications of the probabilistic voting model regarding the behavior of indifferent voters.

5.2 Win Motivation

Under win motivation in the stochastic preference model, we again have discontinuities along the diagonal, posing potential difficulties for equilibrium existence. Despite the presence of these discontinuities, there is still a unique equilibrium when the policy space is unidimensional, as with vote motivation, though now the characterization is changed. Let H_μ denote the distribution of median ideal policy, i.e., $H_\mu(x)$ is the probability that the median voter's ideal policy is less than or equal to x. By our dispersion assumption, H_μ is strictly increasing and has a unique median, denoted x_μ. The next result, due to Calvert (1985), establishes that in equilibrium the candidates must locate at the same policy, the median of x_μ. Thus, in the unidimensional version of the stochastic preference model, we are back to a median-like result, but now the equilibrium is at the median of the distribution of medians.

Theorem 14 (Calvert) *In the stochastic preference model, assume* X *is unidimensional and win motivation. There is a unique equilibrium* (x_A^*, x_B^*)*. In equilibrium, the candidates locate at the median of the distribution of median ideal policies:* $x_A^* = x_B^* = x_\mu$.

It is relatively easy to show that it is an equilibrium for both candidates to locate at the median of medians x_μ. The uniqueness argument proceeds along the lines of the proof of Theorem 12. If both candidates locate at policy $x^* < x_\mu$, for example, they each win with probability one half. But if candidate A deviates slightly to the right to $x^* + \epsilon < x_\mu$, then with probability greater than one half the median voter will be to the right of $x^* + \epsilon$ and strictly prefer candidate A. By single-peakedness of the voters' utility functions, in such cases candidate A will win a majority of votes, and so A's probability of winning after deviating is greater than one half, a contradiction. As with vote motivation, it is clear that the equilibrium policies in the stochastic preference model converge to the median ideal policy when we consider models approaching deterministic voting in the Downsian model.

For multidimensional policy spaces, the equilibrium result in Theorem 13 for the representative voter model carries over directly, for in this model the objectives of vote motivation and win motivation coincide. Thus, similar negative conclusions hold for equilibrium existence in multiple dimensions. As with vote motivation, however, Duggan and Jackson (2005) prove existence of mixed strategy equilibria.

5.3 Policy Motivation

In the general probabilistic voting framework, policy motivation for the candidates takes the form

$$EU_A(x_A, x_B) \quad = \quad P(x_A, x_B)u_A(x_A) + (1 - P(x_A, x_B))u_A(x_B),$$

and likewise for candidate B. We maintain the focus of our analysis on the stochastic preference model, except for our first result, which gives an easy necessary condition for equilibria that holds for both probabilistic voting models. The result, which is proved by Wittman (1983, 1990),[15] Hansson and Stuart (1984), Calvert (1985), and Roemer (1994), shows that the candidates can never locate at identical positions in equilibrium.

Theorem 15 (Wittman; Hansson and Stuart; Calvert; Roemer) *In the stochastic partisanship or stochastic preference models, assume policy motivation. If (x_A^*, x_B^*) is an equilibrium, then the candidates do not locate at the same policy position: $x_A^* \neq x_B^*$.*

To prove this result, simply note that whenever the two candidates locate at the same position, at least one candidate can deviate profitably by moving

[15] Wittman's (1983, 1990) model is slightly different from the one here, as he assumes a hybrid of vote and policy motivation.

toward his or her ideal policy. By our assumption on the support of F_i in the stochastic partisanship model, and by our dispersion assumption in the stochastic preference model, that candidate will win with positive probability after moving, increasing his or her expected utility.

By Theorem 15, if we consider a stochastic preference model close to the Downsian model, then the candidates will adopt distinct policies in equilibrium. Assuming a unidimensional policy space, the median voter theorem implies that, in the Downsian model, the unique equilibrium is for both candidates to locate at the median ideal policy. Calvert (1985) and Roemer (1994) show that, as the stochastic preference model gets close to Downsian, the wedge between the candidates' equilibrium policies goes to zero, and the equilibrium policies of both candidates', if equilibria exist, must converge to the median ideal policy. Thus, equilibria appear to change in a continuous way when we add noise to the behavior of voters in the Downsian model, as we have seen under vote and win motivation.

Theorem 15 leaves open the question of whether an equilibrium exists in the first place. Under our maintained assumptions, measure-theoretic arguments can be used to show that each voter i's probability $P_i(x_A, x_B)$ is continuous in the positions of the candidates whenever x_A and x_B are distinct, and, as discontinuities no longer occur when one candidate crosses over another, the candidates' objective functions under policy motivation are therefore continuous. Thus, one of the main conditions in the DFG theorem is fulfilled. But the candidates' objective functions may not be quasi-concave, making equilibrium existence a non-trivial issue. Duggan and Fey (2005a) show that if the distribution of vectors $(\theta_1, \ldots, \theta_n)$ piles probability mass near a particular vector of parameters, and if the core is empty for those parameters, then there does not exist an equilibrium in the probabilistic voting model. This result is relevant when the policy space is multidimensional and candidates have a large amount of information about voter preferences.

In the quadratic stochastic preference model with a unidimensional policy space, assuming the distribution G of ideal policies is uniform over an interval containing $[\tilde{x}_A, \tilde{x}_B]$, the candidates' objective functions are strictly concave over $[\tilde{x}_A, \tilde{x}_B]$. This allows us to apply the DFG theorem, a simple observation noted by Duggan and Fey (2005b). Though the result is stated only for a uniform distribution, it holds as well for distributions close enough to uniform.

Theorem 16 (Duggan and Fey) *In the quadratic stochastic preference model, assume X is unidimensional, policy motivation, and G is uniform on an interval containing $[\tilde{x}_A, \tilde{x}_B]$. There is an equilibrium.*

In case u_A and u_B are quadratic and G is uniform with mean at the midpoint of $\tilde{x}_A$ and $\tilde{x}_B$, there is a unique equilibrium, and we can calculate it explicitly.

Assume for simplicity that G is uniform on $[\tilde{x}_A, \tilde{x}_B] = [0, 1]$. The first order condition for candidate A reduces to the simple quadratic equation,

$$-3x_A^2 - 2x_B x_A + x_B^2 = 0,$$

with the solution $x_A = x_B/3$. This is indeed a best response to x_B, and a similar calculation for candidate B yields $x_B = (x_A/3) + (2/3)$. Solving these equations yields $x_A = 1/4$ and $x_B = 3/4$, giving us the unique equilibrium for this special case.

What if we depart from the assumption of uniformly distributed voter ideal policies in Theorem 16? Perhaps surprisingly, equilibria can fail to exist due to convexity problems, even in the unidimensional model. Duggan and Fey (2005b) show that this failure of existence can occur even in highly structured settings — with quadratic utilities for the candidates and a symmetric, single-peaked density for the distribution of ideal policies.[16] They assume quadratic utilities with $\tilde{x}_A = 0$ and $\tilde{x}_B = 1$ and use a piece-wise linear distribution G that puts probability arbitrarily close to one near 1/2. The problem created by the lack of quasi-concavity is that the candidates' best responses are not uniquely defined, creating the possibility of a "jump" in their reaction functions. In Figure 6, this occurs for candidate A at approximately $x_B = .6$. In response to this policy position, A is indifferent between adopting a relatively desirable policy (about .2), but winning with a lower probability, and adopting a less appealing policy (about .4), but winning with a higher probability. This is also true of candidate B. As a result, in this example, the reaction functions of the candidates do not cross, and an equilibrium fails to exist.[17]

As a consequence, we see that the continuity result of Calvert (1985) and Roemer (1994) can be vacuous for some stochastic preference models: it may be that, when we add arbitrarily small amounts of noise to voting behavior in the Downsian model, pure strategy equilibria simply do not exist. Because the objective functions of the candidates are continuous in the stochastic preference model under policy motivation, however, the DFG theorem yields a general existence result for mixed strategy equilibrium, regardless of the dimensionality of the policy space.

Theorem 17 *In the stochastic preference model, assume policy motivation. There exists a mixed strategy equilibrium.*

[16] Ball (1999) presents an example of equilibrium non-existence in the model of mixed motivations. His example exploits the discontinuity in that model introduced by a positive weight on holding office.

[17] Hansson and Stuart (1984) claim that the an equilibrium exists if each candidate's probability of winning is concave in the candidate's own position, but their claim rests on the incorrect assumption that the candidates' objective functions are then concave. Thus, the question of general sufficient conditions for existence is open.

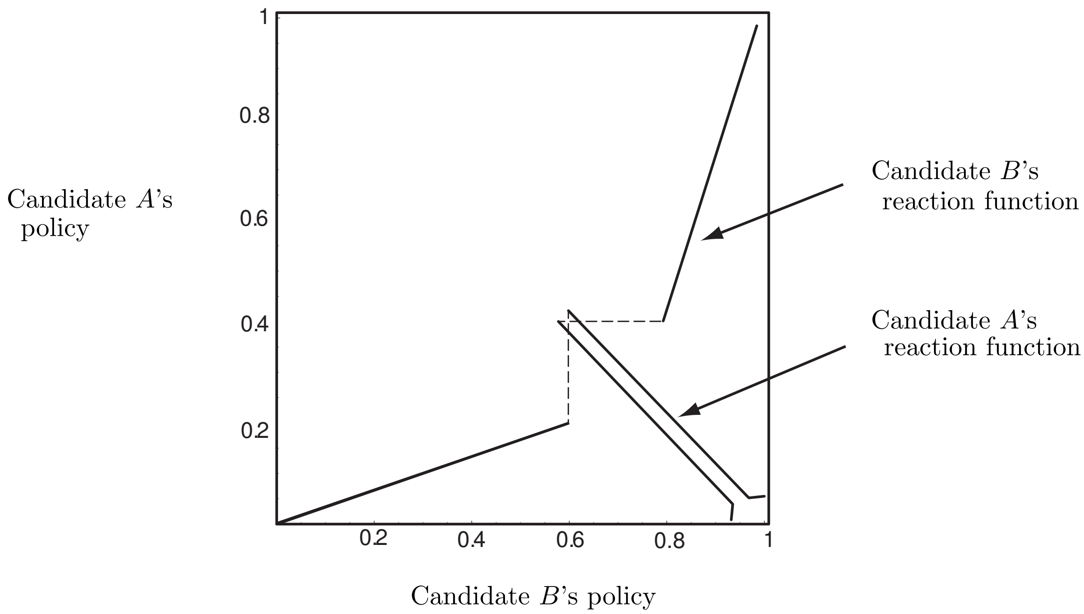

Figure 6: Equilibrium non-existence with policy motivation

Finally, consider a stochastic preference models close to the Downsian model. Assuming a unidimensional policy space, Theorem 3 implies that in the Downsian model the unique equilibrium under policy motivation is for both candidates to locate at the median ideal policy. As we add a small amount of noise to the behavior of voters in the Downsian model, we have seen that pure strategy equilibria can fail to exist. By Theorem 17, there will be mixed strategy equilibria, however, and Duggan and Jackson (2005) show that, as the model gets close to Downsian, the mixed strategy equilibria must put probability arbitrarily close to one on policies close to the median ideal policy. Thus, equilibria change in a continuous way when we add noise to the Downsian model.

6 Conclusion

Of many themes throughout this article, the most prominent has been the difficulty in ensuring existence of equilibria. This is especially true for the Downsian model when the policy space is multidimensional. Probabilistic voting models eliminate some of the discontinuities of the Downsian model, and the analysis of these models has yielded reasonable (if somewhat restrictive) sufficient conditions for equilibrium existence under vote motivation, ensuring that the candidates' objective functions have the appropriate convexity properties. These conditions do not always hold under win motivation or policy motivation or when voting behavior is close to Downsian. As a consequence equilibria may fail to exist, even when the policy space is unidimensional.

I have pursued one approach to solving the existence problem, namely, analyzing mixed strategy equilibria. We have seen that these exist very generally, and that they often possess continuity properties that are desirable in a formal modelling approach: when we add a small amount of noise to voting behavior in the Downsian model, the equilibrium mixed strategies of the candidates are close to the median ideal policy. This assures us that the predictions of the Downsian median voter theorem are not unduly sensitive to the specifications of the model.

Other approaches have been pursued in the literature, though I do not cover them at length in this article. One approach, taken by Roemer (2001), is to modify the objectives of the candidates to demand more of a deviation to be profitable. Roemer endows parties (rather than individual candidates) with multiple objectives of office and policy motivation, as well as an interest in "publicity," whereby a party seeks to announce policy platforms consonant with its general stance, regardless of whether these platforms win. Roemer shows that, if we assume a deviation must satisfy all three of these objectives to be profitable, then an equilibrium exists in some two-dimensional environments.

A second approach, referred to as "citizen candidate" models, is pursued by Osborne and Slivinski (1996) and Besley and Coate (1997). They model candidates as policy motivated and assume candidates cannot commit to campaign promises, removing all positional aspects from the electoral model. Instead, the strategic variable is whether to run in the election, and equilibria are guaranteed to exist. A related approach, referred to as "electoral accountability" models, views elections as repeated over time and takes up informational aspects of elections. Work in this vein, such as Ferejohn (1986), Banks and Sundaram (1993, 1998), Duggan (2000b), and Banks and Duggan (2002), again abstracts away from campaigns. Politicians do make meaningful choices while in office, however, as candidates must consider the information (about preferences or abilities) their choices convey to voters. Equilibria of a simple (stationary) form can often be shown to exist quite generally in these models.

References

[1] P. Aranson, M. Hinich, and P. Ordeshook (1974) "Election Goals and Strategies: Equivalent and Non-Equivalent Candidate Objectives," *American Political Science Review,* 68: 135–152.

[2] D. Austen-Smith and J. Banks (1999) *Positive Political Theory I: Collective Preference,* Ann Arbor: University of Michigan Press.

[3] D. Austen-Smith and J. Banks (2004) *Positive Political Theory II: Strategy and Structure,* forthcoming, Ann Arbor: University of Michigan Press.

[4] R. Ball (1999) "Discontinuity and Non-Existence of Equilibrium in the Probabilistic Spatial Voting Model," *Social Choice and Welfare,* 16: 533–556.

[5] J. Banks (1995) "Singularity Theory and Core Existence in the Spatial Model," *Journal of Mathematical Economics,* 24: 523–536.

[6] J. Banks and J. Duggan (2002) "A Multidimensional Model of Repeated Elections," mimeo.

[7] J. Banks and J. Duggan (2005) "Probabilistic Voting in the Spatial Model of Elections: The Theory of Office-Motivated Candidates," in *Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks,* D. Austen-Smith and J. Duggan, eds., New York: Springer-Verlag.

[8] J. Banks and R. Sundaram (1993) "Moral Hazard and Adverse Selection in a Model of Repeated Elections," in *Political Economy: Institutions, Information, Competition and Representation,* W. Barnett et al. eds., New York: Cambridge University Press.

[9] J. Banks and R. Sundaram (1998) "Optimal Retention in Agency Problems," *Journal of Economic Theory,* 82: 293–323.

[10] T. Besley and S. Coate (1997) "An Economic Model of Representative Democracy," *Quarterly Journal of Economics,* 112: 85–114.

[11] D. Black (1958) *The Theory of Committees and Elections,* Cambridge: Cambridge University Press.

[12] R. Calvert (1985) "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence," *American Journal of Political Science,* 29: 69–95.

[13] R. Calvert (1986) *Models of Imperfect Information in Politics,* London: Harwood Press.

[14] P. Coughlin (1990) "Candidate Uncertainty and Electoral Equilibria," *Advances in the Spatial Theory of Voting,* J. Enelow and M. Hinich, eds., Cambridge: Nwe York.

[15] P. Coughlin (1992) *Probabilistic Voting Theory,* Cambridge: Cambridge University Press.

[16] P. Coughlin and S. Nitzan (1981) "Electoral Ourcomes with Probabilistic Voting and Nash Social Welfare Maxima," *Journal of Public Economics,* 15: 113–121.

[17] G. Cox (1984) "Non-Collegial Simple Games and the Nowhere Denseness of the Set of Preference Profiles having a Core," *Social Choice and Welfare*, 1: 159–164.

[18] O. Davis, M. Hinich, and P. Ordeshook (1970) "An Expository Development of a Mathematical Model of the Electoral Process," *American Political Science Review,* 64: 426–448.

[19] G. Debreu (1952) "Existence of a Social Equilibrium," *Proceedings of the National Academy of Sciences,* 38: 886–893.

[20] A. Downs (1957) *An Economic Theory of Democracy,* New York: Harper and Row.

[21] J. Duggan (2000a) "Equilibrium Equivalence Under Expected Plurality and Probability of Winning Maximization," mimeo.

[22] J. Duggan (2000b) "Repeated Elections with Asymmetric Information," *Economics and Politics,* 12: 109–136.

[23] J. Duggan (2005) "A Note on Uniqueness of Electoral Equilibrium When the Median Voter is Unobserved," mimeo.

[24] J. Duggan and M. Fey (2005a) "Electoral Competition with Policy-Motivated Candidates," *Games and Economic Behavior,* 51: 490–522.

[25] J. Duggan and M. Fey (2005b) "Indeterminacy in Elections with Policy-Motivated Candidates and a Single Issue Dimension," mimeo.

[26] J. Duggan and M. Jackson (2005) "Mixed Strategy Equilibrium and Deep Covering in Multidimensional Electoral Competition," mimeo.

[27] K. Fan (1952) "Fixed-Point and Minimax Theorems in Locally Convex Linear Spaces," *Proceedings of the National Academy of Sciences,* 38: 121–126.

[28] J. Ferejohn (1986) "Incumbent Performance and Electoral Control," *Public Choice,* 50: 5–25.

[29] I. Glicksberg (1952) "A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium," *Proceedings of the American Mathematical Society,* 3: 170–174.

[30] I. Hansson and C. Stuart (1984) "Voting Competitions with Interested Politicians: Platforms Do Not Converge to the Preferences of the Median Voter," *Public Choice,* 44:431–441.

[31] M. Hinich (1977) "Equilibrium in Spatial Voting: The Median Voter Result is an Artifact," *Journal of Economic Theory* 16: 208–219.

[32] M. Hinich (1978) "The Mean Versus the Median in Spatial Voting Games," in P. Ordeshook, ed., *Game Theory and Political Science,* New York: NYU Press.

[33] M. Hinich, J. Ledyard, and P. Ordeshook (1972) "Nonvoting and the Existence of Equilibrium Under Majority Rule," *Journal of Economic Theory,* 4: 144–153.

[34] M. Hinich, J. Ledyard, and P. Ordeshook (1973) "A Theory of Electoral Equilibrium: A Spatial Analysis Based on the Theory of Games," *Journal of Politics,* 35: 154–193.

[35] H. Hotelling (1929) "Stability in Competition," *Economic Journal,* 39: 41–57.

[36] R. Hoyer and L. Mayer (1974) "Comparing Strategies in a Spatial Model of Electoral Competition," *American Journal of Political Science,* 18: 501–523.

[37] G. Kramer (1978) "Robustness of the Median Voter Result," *Journal of Economic Theory,* 19: 565–567.

[38] D. Laussel and M. Le Breton (2002) "Unidimensional Downsian Politics: Median, Utilitarian, or What Else?" *Economics Letters,* 76: 351–356.

[39] J. Ledyard (1984) "The Pure Theory of Large Two-Candidate Elections," *Public Choice,* 44: 7–41.

[40] A. Lindbeck and J. Weibull (1987) "Balanced-Budget Redistribution as the Outcome of Political Competition," *Public Choice,* 52: 273–297.

[41] A. Lindbeck and J. Weibull (1993) "A Model of Political Equilibrium in a Representative Democracy," *Journal of Public Economics,* 51: 195–209.

[42] R. McKelvey (1986) "Covering, Dominance, and Institution-free Properties of Social Choice," *American Journal of Political Science,* 30: 283–314.

[43] J. Nash (1950) "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences,* 36: 48–49.

[44] P. Ordeshook (1986) *Game Theory and Political Theory,* New York: Cambridge Press.

[45] M. Osborne (1995) "Spatial Models of Political Competition Under Plurality Rule: A Survey of Some Explanations of the Number of Candidates and the Positions They Take," *Canadian Journal of Economics,* 28: 261–301.

[46] M. Osborne and A. Slivinski (1996) "A Model of Political Competition with Citizen-Candidates," *Quarterly Journal of Economics,* 111: 65–96.

[47] J. Patty (2003) "Local Equilibrium Equivalence in Probabilistic Voting Models," mimeo.

[48] C. Plott (1967) "A Notion of Equilibrium and Its Possibility Under Majority Rule," *American Economic Review,* 57: 787–806.

[49] W. Riker (1980) "Implications from the Disequilibrium of Majority Rule for the Study of Institutions," *American Political Science Review,* 74: 432–446.

[50] J. Roemer (1994) "A Theory of Policy Differentiation in Single Issue Electoral Politics," *Social Choice and Welfare,* 11: 355–380.

[51] J. Roemer (2001) *Political Competition: Theory and Applications,* Cambridge: Harvard University Press.

[52] D. Saari (1997) "The Generic Existence of a Core for q-Rules," *Economic Theory,* 9: 219–260.

[53] K. Shepsle (1991) *Models of Multiparty Electoral Competition,* London: Harwood Press.

[54] D. Wittman (1977) "Candidates with Policy Preferences: A Dynamic Model," *Journal of Economic Theory,* 14: 180–189.

[55] D. Wittman (1983) "Candidate Motivation: A Synthesis of Alternative Theories," *American Political Science Review,* 77: 142–157.

[56] D. Wittman (1990) "Spatial Strategies When Candidates Have Policy Preferences," *Advances in the Spatial Theory of Voting,* J. Enelow and M. Hinich, eds., New York: Cambridge Press.