A Survey of Equilibrium Analysis in Spatial Models of Elections*

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

March 28, 2017

*I thank Jon Eguia, David Epstein, Mark Fey, Cesar Martinelli, and Martin Osborne (on repeated occasions!) for helpful feedback.

Contents

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 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 but suffers from existence failures in multiple dimensions, models of probabilistic voting are commonly thought to mitigate existence problems. We will see that this is true to an extent, but that probabilistic voting can actually introduce existence problems where there were none before — even 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 (or mitigating) 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.

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

[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.

number of surveys, such as Wittman (1990), Coughlin (1990, 1992), Austen-Smith and Banks (2004), and Duggan (2006a). 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. The term *office motivation* is generally used to refer to either of the next two objectives, where candidates are both win-motivated or both are vote-motivated. Our first objective function 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 $(2\sum b_i) - n$.

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 further assume that 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; and that these ideal policies are interior to the policy space, i.e., they lie strictly inside the boundary of X. Then candidate A's utility from policy positions x_A and x_B and ballots $(b_1, \ldots, b_n)$ is

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

and likewise for candidate B.

A third approach combines the latter two objectives.

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.

Voting behavior. In order to model voting behavior, we represent each voter i's preferences by an objective function $U_i(C, x, \theta_i)$ in general. Here, C is the candidate who wins the election, x is the policy platform of the winning candidate C, and θ_i is a preference parameter.[2] The vector $(\theta_1, \ldots, \theta_n)$ of parameters is not directly observed by the candidates, so we model θ_i as a random variable from the candidates' point of view. Note that this general formulation allows for uncertainty on the part of candidates, but a special case is the model in which the distribution of voter i's parameter is degenerate, i.e., it concentrates probability one on a single value of θ_i, in which case the candidates know the voters' preferences with probability one. At this point, we can allow these parameters to be correlated, but we focus on the voters' marginal vote probabilities. In short, we assume a voter votes for the candidate-policy pair (C, x_C) offering the highest utility given parameter θ_i, and flipping a coin when indifferent. 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 strategies in voting subgames.

Letting $P_i(x_A, x_B, \theta_i)$ denote the probability that voter i votes for candidate A, i.e., casts ballot $b_i = 1$, given parameter θ_i, we then have

$$P_i(x_A, x_B, \theta_i) = \begin{cases} 1 & \text{if } U_i(A, x_A, \theta_i) > U_i(B, x_B, \theta_i) \\ 0 & \text{if } U_i(A, x_A, \theta_i) < U_i(B, x_B, \theta_i) \\ \frac{1}{2} & \text{else,} \end{cases}$$

and the probability i votes for B is $1 - P_i(x_A, x_B, \theta_i)$. From the point of view of the candidates, who do not observe $(\theta_1, \ldots, \theta_n)$, the probability that voter i casts her ballot in favor of A is obtained by "integrating out" the parameter, and we write this as

$$\begin{aligned}P_i(x_A, x_B) &= \Pr(\theta_i \mid U_i(A, x_A, \theta_i) > U_i(B, x_B, \theta_i))\\&\quad + \frac{1}{2}\Pr(\theta_i \mid U_i(A, x_A, \theta_i) = U_i(B, x_B, \theta_i)).\end{aligned}$$

Typically, the parameter is continuously distributed, and (when the candidates choose distinct platforms) the probability of exact indifference is zero, so the last term in the above expression drops out.

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

[2] Implicit in this representation is the assumption that voters do not care about a candidate's margin of victory or the platform of the losing candidate.

is taken over vectors of ballots $(b_1, \ldots, b_n)$ with respect to the distribution induced by the parameters $(\theta_1, \ldots, \theta_n)$, and likewise for candidate B. Since 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 Downsian Voting

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 utility function $u_i(x, \theta_i)$ that is strictly concave and differentiable in x (so $u_i(\cdot, \theta_i)$ satisfies the conditions imposed on candidate utilities under policy motivation). Then voter i's objective function is $U_i(C, x, \theta_i) = u_i(x_C, \theta_i)$, i.e., the voter cares only about the policy position of the winning candidate. To capture the idea that candidates know voter preferences with certainty, we assume that the distribution of the parameter θ_i is degenerate; for simplicity, we assume that θ_i is a real number, and that it equals zero with probability one for all voters. Thus, we write $u_i(x)$ for the policy utility of voter i instead of $u_i(x, 0)$, and the probability that voter i casts her ballot for candidate A is then

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

so only the ballots of indifferent voters (who randomize with equal probability) are unpredictable by the candidates.

Under our maintained assumptions, voter i has a unique *ideal policy,* denoted $\tilde{x}_i$, which uniquely maximizes u_i and is strictly preferred by the voter to 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, i.e., for all distinct i and j, we assume $\tilde{x}_i \neq \tilde{x}_j$.

3.1 Social Choice Background

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

$$x\ M\ y \quad \Leftrightarrow \quad \#\{i \mid u_i(x) > u_i(y)\}\ >\ \frac{n}{2}.$$

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.

In fact, the above observations follow from a general property of majority preferences. We first note that even though voter utilities are strictly concave, majority preferences do not inherit the full convexity properties of voter preferences: it can easily be the case that two policies, say y and z, are majority preferred to a policy x, but some policy between y and z is not. The reason this is possible is because y and z may be preferred to x by *different* majorities. (See Figure 3 for a typical example of this.) But majority preferences do possess the technical-sounding property of *strict-starshapedness,* which means that for any two policies x and y, if it is not the case that y is majority-preferred to x, then every policy between x and y is majority-preferred to y. Formally, if for two distinct policies it is not the case that $y\ M\ x$, then for all α strictly between zero and one, we have $z = \alpha y + (1-\alpha)x\ M\ y$. Although strange at first sight, this observation is useful in the analysis of some kinds of electoral models.

Lemma 1 *Majority preferences M are strictly star-shaped.*

The result is straightforward to prove. If it is not the case that $y\ M\ x$, then

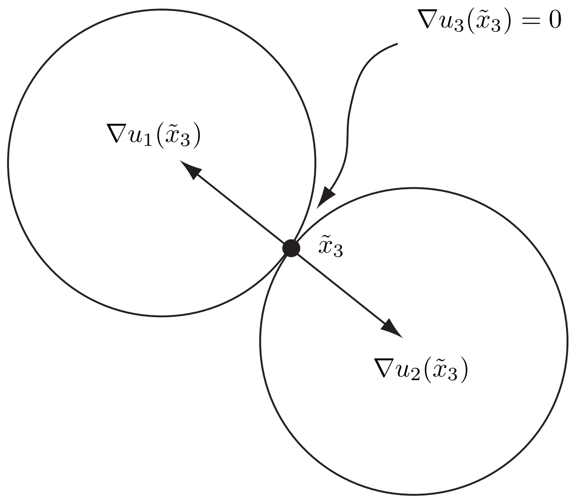

Figure 1: Radial symmetry

x must offer at least half of the voters at least as much utility as y, i.e.,

$$\#\{i \in N \mid u_i(x) \geq u_i(y)\} \quad \geq \quad \frac{n}{2},$$

and since n is odd the latter inequality must hold strictly. Letting $z = \alpha x + (1 - \alpha) y$ be any policy strictly between x and y, strict concavity (in fact, strict quasi-concavity is enough) of voter utilities implies that $u_i(z) > u_i(y)$ for every voter with $u_i(x) \geq u_i(y)$. Since these voters comprise a majority, we have $z \; M \; y$, as required.

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*

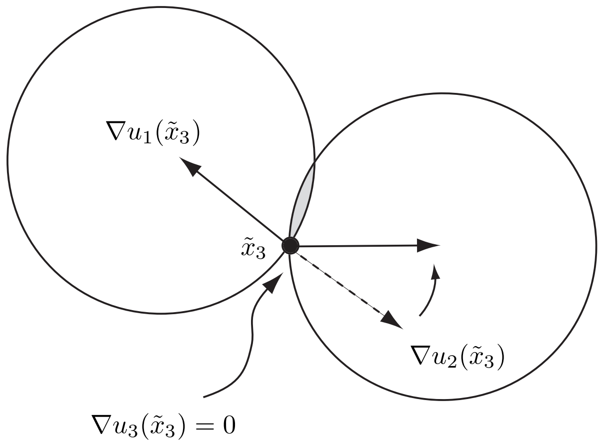

Figure 2: Violation of radial symmetry

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. 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.[3]

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

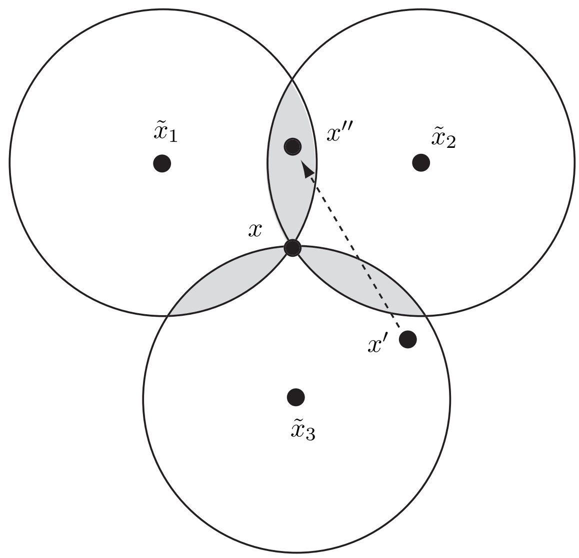

Figure 3: Discontinuity and non-convexity

3.2 Office Motivation

Under the assumption of deterministic voting, a candidate wins for sure if her policy platform is majority-preferred to the others; she loses for sure when the situation is reversed; and in the remaining case, where neither candidate's policy is majority-preferred to the others, they each win (and lose) with positive probability. Thus, win motivation satisfies the following:

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

and otherwise $0 < EU_A(x_A, x_B) < 1$; 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.

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 vote 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 office-motivation; and assume* X *is unidimensional. 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.[4]

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.

3.3 Policy Motivation

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

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

and otherwise, if neither candidate's platform is majority-preferred to the other's, then we have $u_A(x_A) > EU_A(x_A, x_B) > u_A(x_B)$ when $u_A(x_A) > u_A(x_B)$, and we have $EU_A(x_A, x_B) = u_A(x_A) = u_A(x_B)$ when the candidate is indifferent between the two platforms. The above objective function also suffers from discontinuities and non-convexities, but it differs fundamentally from the case of

[4]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.

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 policy motivation; and X is unidimensional. 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 (2005) 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 (2005) show that it is a near universal feature of electoral competition with policy motivated candidates, regardless of the dimensionality of the policy space. Roemer (2001) also argues that equilibria in which the two candidates adopt distinct policy positions almost never exist. The proof, though composed of very simple steps, is the most involved of those provided in this survey; I include it for the interested reader.

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., $x_A^* \neq \tilde{x}_A$ and $x_B^* \neq \tilde{x}_B$, then the candidates' policy positions are identical: $x_A^* = x_B^* = x^*$.*

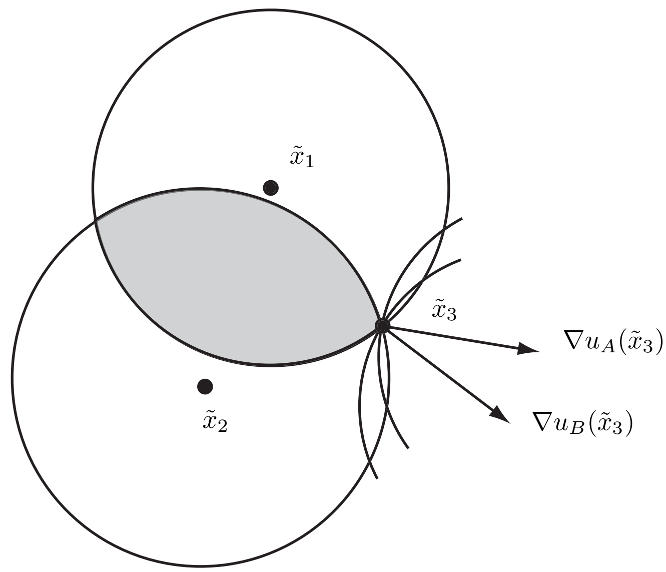

Figure 4: Equilibrium with policy motivation but not win motivation

To see this result, suppose in order to deduce a contradiction that (x_A^*, x_B^*) is an equilibrium in which neither candidate locates at her ideal policy, but $x_A^* \neq x_B^*$. Since it cannot be that the two policies are majority-preferred to each other, we assume without loss of generality that it is not the case that x_B^* is majority-preferred to x_A^*, i.e., not $x_B^* \; M \; x_A^*$. I claim that $u_A(x_A^*) > u_A(x_B^*)$, for otherwise $u_A(x_B^*) \geq u_A(x_A^*)$. We cannot have $u_A(x_B^*) > u_A(x_A^*)$, because then candidate A could deviate to B's platform and obtain a higher level of expected utility, i.e., $EU_A(z, x_B^*) = u_A(x_B^*) > EU_A(x_A^*, x_B^*)$, but this contradicts the fact that (x_A^*, x_B^*) is an equilibrium. Thus, it must be that $u_A(x_A^*) = u_A(x_B^*)$, and then strict concavity (in fact, strict quasi-concavity) of u_A implies that candidate A strictly prefers every policy between the two platforms to both x_A^* and x_B^*. In particular, letting $z = \frac{1}{2}x_A^* + \frac{1}{2}x_B^*$ be the midpoint between the platforms, we have $u_A(z) > u_A(x_A^*) = u_A(x_B^*)$. But by strict starshapedness, from Lemma 1, it follows that z is majority-preferred to x_B^*. Then we have

$$EU_A(z, x_B^*) \; = \; u_A(z) \; > \; EU_A(x_A^*, x_B^*),$$

but then moving to z is a profitable deviation for A, contradicting the fact that (x_A^*, x_B^*) is an equilibrium. This establishes the claim. I next claim that it is not the case that x_A^* is majority-preferred to x_B^*, for otherwise there is a majority of voters for whom x_A^* is strictly preferred to x_B^*, i.e., $u_i(x_A^*) > u_i(x_B^*)$. By differentiability (in fact, continuity) of voter utilities, these strict inequalities hold for all policies very close to x_A^*. But then candidate A could move from x_A^* to some policy y slightly in the direction of her gradient, y would still be strictly majority-preferred to x_B^*, and we would have

$$EU_A(y, x_B^*) \;=\; u_A(y) \;>\; u_A(x_A^*) \;\geq\; EU_A(x_A^*, x_B^*),$$

where the strict inequality follows from strict concavity of u_A. But then moving to y is a profitable deviation for A, again a contradiction. We conclude that not $x_A^* \; M \; x_B^*$, and we are prepared for a final application of the star-shapedness lemma. By Lemma 1, every policy between x_A^* and x_B^* is majority-preferred to x_B^*, so consider such a policy v that is slightly in the direction of x_B^* but very close to x_A^*; formally, we write $v = (1-\alpha)x_A^* + \alpha x_B^*$, where α is very close to, but greater than, zero. Recall that since neither candidate's platform is majority-preferred to the other's, we have $u_A(x_A^*) > EU_A(x_A^*, x_B^*)$. By differentiability (in fact, continuity) of u_A, if v is close enough to x_A^* (i.e., α is close enough to zero), then this strict inequality is maintained. Combining these observations, we have

$$EU_A(v, x_B^*) = u_A(v) > EU_A(x_A^*, x_B^*),$$

but then moving to v is a profitable deviation for A, a contradiction. Therefore, we have deduced a contradiction in all possible cases, and we conclude that our initial supposition that $x_A^* \neq x_B^*$ is false, i.e., that $x_A^* = x_B^*$, as required.

The next result shows that the kind of positive situation depicted in Figure 4 is limited to the two-dimensional case. The result, due to Duggan and Fey (2005), 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 x^* is interior to the policy space 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 $j \notin \{i, k\}$ (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.

3.4 Mixed Motivation

Mixed motivations take the following form in the Downsian model:

$$EU_A(x_A, x_B) = \begin{cases} u_A(x_A) + w & \text{if } x_A \; M \; x_B \\ u_A(x_B) & \text{if } x_B \; M \; x_A, \end{cases}$$

and otherwise, if neither candidate's platform is majority-preferred to the other's, then we have $u_A(x_A) + w > EU_A(x_A, x_B)$ when $u_A(x_A) + w > u_A(x_B)$, and we have $EU_A(x_A, x_B) = u_A(x_A) + w = u_A(x_B)$ when the candidate is indifferent between winning and losing.

As in the case of pure office or pure policy motivation, we are back to a median voter result when the policy space is one-dimensional, and existence of pure strategy equilibria in multiple dimensions is problematic. In fact, the predicament is worse than that stated in Theorem 5, as we are now confronted with the situation of Theorem 2: equilibria exist only if the core is nonempty, which typically fails to occur.

Theorem 6 (Duggan and Fey) *In the Downsian model, assume mixed motivation. If there is an equilibrium (x_A^*, x_B^*) such that neither candidate locates at his or her ideal policy, i.e., $x_A^* \neq \tilde{x}_A$ and $x_B^* \neq \tilde{x}_B$, then the core is non-empty. In this case, the equilibrium is unique, and the candidates locate at the core point: $x_A^* = x_B^* = \tilde{x}_k$. Conversely, if the core is non-empty, then the platform pair $(\tilde{x}_k, \tilde{x}_k)$ is an equilibrium.*

Let (x_A^*, x_B^*) be an equilibrium of the electoral game under mixed motives such that neither candidate locates at his or her ideal point. By the arguments in the proof of Theorem 4, the candidates adopt identical platforms, say $x_A^* = x_B^* = x^*$. Note that $EU_A(x_A^*, x_B^*) = u_A(x^*) + \frac{w}{2}$. Suppose, in order to deduce a contradiction, that x^* does not belong to the core, so candidate A can move to a policy z that is majority-preferred to x^*. Though this move may not be appealing to A on policy grounds, consider a move by A from x^* to y slightly in the direction of z; formally, we define $y = (1 - \alpha)x^* + \alpha z$ with α close to, but greater than, zero. By strict star-shapedness, from Lemma 1, it follows that y is majority preferred to x^*. Furthermore, as α approaches zero, candidate A's expected utility from y converges to $u_A(x^*) + w$; formally, we take the limit as α goes to zero, and continuity of u_A implies that

$$\lim_{\alpha \to 0} EU_A(y, x_B^*) \; = \; \lim_{\alpha \to 0} u_A(y, x_B^*) + w \; = \; u_A(x^*) + w$$

$$> \quad u_A(x^*) + \frac{w}{2} \;=\; EU_A(x_A^*, x_B^*).$$

But this is higher than A's utility from the equilibrium policy x^*, contradicting the fact that (x_A^*, x_B^*) is an equilibrium. We conclude that x^* belongs to the core, as required.

The proof of Theorem 6 is considerably simpler than for Theorem 5 (which is omitted). At work here is a discontinuity introduced by the reward to winning: as A chooses majority-preferred platforms close to x^*, the policy difference between x^* and the deviation become negligible, but the positive reward w does not. Therefore, we conclude that x^* must be a core point. But we know from Plott's theorem that such points almost never exist, as claimed. As previously, the result of Duggan and Jackson (2005) yields a mixed strategy equilibrium in the extended Downsian model.

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. This is the approach taken here.

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, who view these preference parameters as random variables.

Stochastic partisanship model. Assume that each voter i has a strictly concave, differentiable utility function u_i with unique ideal point $\tilde{x}_i$, as in the Downsian voting model. Furthermore, θ_i is a one-dimensional parameter reflecting a "utility bias" θ_i in favor of candidate B.[5] Then voter i's objective function is

$$U_i(C, x, \theta_i) = \begin{cases} u_i(x) & \text{if } C = A \\ u_i(x) + \theta_i & \text{if } C = B, \end{cases}$$

and i votes for A if and only if the utility of candidate A's platform exceeds that of B's platform by at least θ_i. That is, i votes for A if and only if $u_i(x_A) \geq u_i(x_B) + \theta_i$, i.e., $\theta_i \leq u_i(x_A) - u_i(x_B)$. We assume that the profile $(\theta_1, \ldots, \theta_n)$ of biases is a random variable from the candidates' perspective, and we assume that each θ_i is distributed according to the distribution F_i. (For now, the biases are not necessarily independently distributed.) We assume that each F_i is differentiable on an interval that includes all possible utility differences $u_i(x) - u_i(y)$, as x and y range over all of X, with a density f_i that is positive on that interval.[6] implying that the probability that θ_i exactly equals a given utility difference is zero. Then the probability that voter i votes for candidate A given platforms x_A and x_B is $P_i(x_A, x_B) = F_i(u_i(x_A) - u_i(x_B)) \in (0, 1)$. We do not assume the parameters θ_i are distributed symmetrically around zero, nor do we assume they are distributed symmetrically across voters. Convenient special cases include uniform (with sufficiently large support), logistic, and normal distributions.

We first consider the known results on equilibrium existence and characterization under the assumption of vote-motivated candidates; we then examine the results for win motivation; and we end with a brief analysis of policy and mixed motivation.

4.1 Vote Motivation

In the general probabilistic voting framework, the objective function of vote motivated candidates can be written

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

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.

[5] In case θ_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.

[6] 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.

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

$$\overline{x} \quad = \quad \arg\max_{x\in X} \sum_i f_i(0)u_i(x),$$

where utilities are weighted by the densities of the voters' biases at zero; when biases are distributed symmetrically across voters, so $f_i(0) = f_j(0)$ for all i, j, these weights become irrelevant and can be dropped. This policy is often referred to as the "utilitarian optimum," though this term suggests welfare connotations that are difficult to justify.[7] 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, weighted by the bias densities.

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).[8] The general statement here is due to Banks and Duggan (2005).[9] When voter utilities are quadratic, an implication is Hinich's "mean voter" theorem.

Theorem 7 (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^*). Candidate A's maximization problem, given B's platform, is then

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

and since x_A^* is a best response for candidate A, it satisfies the necessary first order condition,

$$\sum_i f_i(u_i(x_A^*) - u_i(x_B^*))\nabla u_i(x_A^*) \quad = \quad 0. \tag{1}$$

[7]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.

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

[9]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.

Now consider the sum of voter utilities

$$W(x) \quad = \quad \sum_i f_i(u_i(x_A^*) - u_i(x_B^*))u_i(x),$$

with weights equal to the bias densities evaluated at the utility difference for each voter. Viewing these weights as fixed, W is the sum of strictly concave functions of x (multiplied by strictly positive weights), and so W is strictly concave. Furthermore, the first order condition for maximizing $W(x)$ is exactly the expression in (1), so we see that x_A^* solves the first order condition for W. Since W is strictly concave, we conclude that x_A^* is the unique maximizer of W. As well, x_B^* satisfies the necessary first order condition for candidate B,

$$\sum_i f_i(u_i(x_A^*) - u_i(x_B^*))\nabla u_i(x_B^*) \quad = \quad 0,$$

and again we see that x_B^* satisfies the first order condition for W, so x_B^* is the unique maximizer of W. In particular, $x_A^* = x_B^*$, which implies that the weights in the definition of W are equal to $f_i(0)$. Finally, we conclude that $x_A^* = x_B^* = \overline{x}$, as required.

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 fulfilled, for example, if the bias terms of the voters are uniformly distributed over the range of utility differences, 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 8 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.[10] By Theorem 7, if there is an equilibrium, then it is unique, and both candidates locate at the utilitarian point.

Theorem 8 (Hinich, Ledyard, and Ordeshook; Lindbeck and Weibull) *In the stochastic partisanship model, assume vote motivation; and assume the following for each voter i and all policies y:*

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

There exists an equilibrium.

[10]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.

The proof of Theorem 8 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.

The previous theorem has an intuitive geometric meaning, which can be seen with the help of the expected vote function Π_V, which takes n probabilities as arguments:

$$\Pi_V(p_1, \ldots, p_n) \quad = \quad \sum_i p_i.$$

Of course, the candidates' objective functions can then be written as the composition

$$EU_A(x_A, x_B) \quad = \quad \Pi_V(F_1(u_1(x_A) - u_1(x_B)), \ldots, F_n(u_n(x_A) - u_n(x_B))),$$

and likewise for candidate B. The expected vote function, as the sum of individual vote probabilities, is a simple linear function. Given a level of expected vote, the set of probability vectors that correspond to that level, say $\Pi_V(p_1, \ldots, p_n) = c$, is "flat" or "linear." When there are just two voters, the set is a straight line and can be depicted in the Cartesian plane. In Figure 5, each line corresponds to a level of expected vote, with lines further to the northeast corresponding to higher levels. In equilibrium, candidate A seeks to maximize $EU_A(x_A, x_B)$ with respect to x_A, taking x_B as fixed, but we can equivalently view this as a constrained optimization problem using the expected vote function defined above. Let $\mathcal{P}(x_B)$ be the set of probability vectors that candidate A can achieve by varying her platform, i.e.,

$$\mathcal{P}(x_B) \quad = \quad \big\{(F_1(u_1(x) - u_1(x_B)), \ldots, F_n(u_n(x) - u_n(x_B)) \mid x \in X\big\}.$$

This set is the shaded area in Figure 5. Then (x_A^*, x_B^*) is an equilibrium if and only if x_A^* delivers the highest achievable level of expected vote for A, i.e., the vector $(F_1(u_1(x_A^*) - u_1(x_B^*)), \ldots, F_n(u_n(x_A^*) - u_n(x_B^*))$ solves the constrained maximization problem

$$\max_{p_1,\ldots,p_n} \Pi_V(p_1, \ldots, p_n)$$

$$\text{s.t. } (p_1, \ldots, p_n) \in \mathcal{P}(x_B^*),$$

and likewise for candidate B.

In general, the set of achievable probability vectors may not possess nice structure. But under the assumptions of Theorem 8, it is essentially convex.[11]

[11] Specifically, given two probability vectors $p = (p_1, \ldots, p_n)$ and $p' = (p'_1, \ldots, p'_n)$ and any probability vector between them, say $\tilde{p} = (\tilde{p}_1, \ldots, \tilde{p}_n)$, there is some achievable probability vector, say $\hat{p} = (\hat{p}_1, \ldots, \hat{p}_n)$, that is weakly greater in each component.

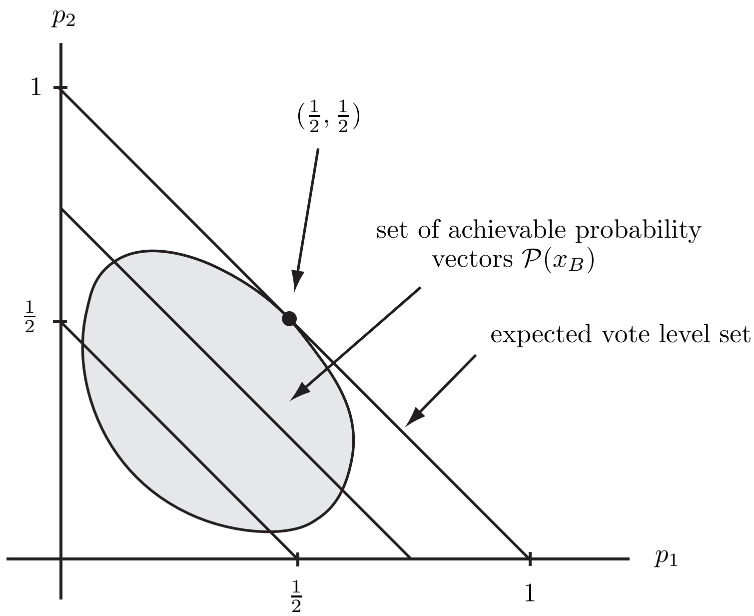

Figure 5: Expected vote maximization

This means that a candidate's best responses, i.e., the probability vectors that attain the highest expected vote level for the candidate given the other's platform, form a convex set, a key property in obtaining existence of an equilibrium. It is especially instructive to let the voters' biases be identically and uniformly distributed over a sufficiently large range; for simplicity, normalize utilities so that the uniform distribution has support on the unit interval. Then $F_i(u_i(x_A) - u_i(x_B)) = u_i(x_A) - u_i(x_B)$, so vote probabilities are just utility differences. When candidate B locates at the utilitarian point, $x_B = \overline{x}$, candidate A's choice of x_A determines a set of utility differences, as in Figure 5, and choosing $x_A = \overline{x}$ determines the vector $(\frac{1}{2}, \ldots, \frac{1}{2})$. Furthermore, by definition $\overline{x}$ solves

$$\max_{x \in X} \sum_i (u_i(x) - u_i(\overline{x})) \quad = \quad \max_{x \in X} \sum_i u_i(x) - c,$$

where $c = \sum_i u_i(\overline{x})$ is a constant from the point of view of A's optimization problem, so $\overline{x}$ is a best response for A, and $\overline{x}$ is an equilibrium. This is reflected in Figure 5, where the set of achievable vote probabilities (i.e., utility differences) for candidate A given $x_B = \overline{x}$ lies below the line corresponding to expected vote level $\frac{n}{2}$.

Theorems 7 and 8 taken together may suggest a puzzling discrepancy between the one-dimensional 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.[12] In this way, we can satisfy the assumptions of the stochastic partisanship model in models arbitrarily close to the Downsian model. By Theorem 7, 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."

In fact, however, Theorem 7 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. Laussel and Le Breton (2002) and Banks and Duggan (2005) 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.[13] 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 9 *In the stochastic partisanship model, assume vote motivation. There exists 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 9 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.

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

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

Stochastic multiplicative partisanship model. Assume each voter i has a differentiable utility function u_i and a utility bias θ_i in favor of candidate B. In contrast to the stochastic partisanship model, we assume that voter utilities and bias terms are positive, that voter utilities are strictly log concave, so $\ln(u_i(x))$ is strictly concave, and that i votes for A if and only if $u_i(x_A) \geq u_i(x_B)\theta_i$. As before, each θ_i is distributed according to a distribution F_i. We assume each F_i is differentiable with density f_i that is positive 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 $F_i(u_i(x_A)/u_i(x_B))$. Let

$$\hat{x} \quad = \quad \arg\max_{x \in X} \prod_i u_i(x)^{f_i(1)}$$

maximize the weighted 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 partisanship 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 7, after an appropriate transformation of the stochastic multiplicative partisanship model. Given a model satisfying the assumptions of the corollary, define an associated stochastic (additive) 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 (additive) partisanship model. Note that $\dot{f}_i(0) = f_i(1)e^0 = f_i(1)$. By Theorem 7, applied to the associated (additive) partisanship model, we therefore have

$$x_A^* \;=\; x_B^* \;=\; \arg\max_{x \in X} \sum_i \dot{f}(0)\dot{u}_i(x) \;=\; \arg\max_{x \in X} \prod_i u_i(x)^{f_i(1)} \;=\; \hat{x},$$

as required.

By a similar logic, we can extend Theorem 8 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 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 7, but any relaxation of concavity of u_i must be offset by greater concavity of F_i in order to apply Theorem 8.

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

- $F_i(\frac{u_i(x)}{u_i(y)})$ *is concave in x*
- $F_i(\frac{u_i(y)}{u_i(x)})$ *is convex in x.*

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 $F_i(u) = \frac{u}{u+1}$. Then by the assumptions of Corollary 3, the function

$$\dot{F}_i(\dot{u}_i(x) - \dot{u}_i(y)) \;=\; \dot{F}_i\left(\ln\left(\frac{u_i(x)}{u_i(y)}\right)\right) \;=\; F_i\left(\frac{u_i(x)}{u_i(y)}\right)$$

is concave in x, and similarly for candidate B. Thus, the assumptions of Theorem 8 are fulfilled for the associated stochastic partisanship model, and we conclude that an equilibrium exists in one model and, therefore, in the other.

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 7 and 8 for the stochastic partisanship model.

4.2 Win Motivation

Finally, let $P(x_A, x_B)$ denote the probability that candidate A wins the election, given the individual vote probabilities. Formally, now assuming the parameters θ_i are independently distributed across voters, this probability of winning has a complicated combinatorial structure:

$$P(x_A, x_B) \quad = \quad \sum_{C:|C|>\frac{n}{2}} \left(\prod_{i\in C} P_i(x_A, x_B)\right)\left(\prod_{i\notin C}(1 - P_i(x_A, x_B))\right).$$

That is, it is just the probability that a majority of voters cast their ballots for A. Of course, B's probability of winning is one minus this amount. Then 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, but lack of quasi-concavity is a problematic issue under win motivation: we will see that even under the concavity conditions of Theorem 8, equilibria may fail to exist in the probability of winning model. Paralleling the analysis of expected vote maximization, it will be useful to define for each vector $(p_1, \ldots, p_n)$ of vote probabilities, the probability of winning function

$$\Pi_W(p_1, \ldots, p_n) \quad = \quad \sum_{C:|C|>\frac{n}{2}} \left(\prod_{i\in C} p_i\right)\left(\prod_{i\notin C}(1 - p_i)\right).$$

Of course, the candidates' objective functions can then be written as the composition

$$EU_A(x_A, x_B) \quad = \quad \Pi_W(F_1(u_1(x_A) - u_1(x_B)), \ldots, F_n(u_n(x_A) - u_n(x_B))),$$

and similarly for candidate B. Then (x_A^*, x_B^*) is an equilibrium if and only if x_A^* delivers the highest achievable probability of winning for A, i.e., the vector $(F_1(u_1(x_A^*) - u_1(x_B^*)), \ldots, F_n(u_n(x_A^*) - u_n(x_B^*))$ solves the constrained maximization problem

$$\max_{p_1,\ldots,p_n} \Pi_W(p_1, \ldots, p_n)$$

$$\text{s.t. } (p_1, \ldots, p_n) \in \mathcal{P}(x_B^*),$$

and likewise for B. See Figure 6.

We know that under the conditions of Theorem 8, locating at the utilitarian point is an equilibrium under vote motivation. Returning to the simple example

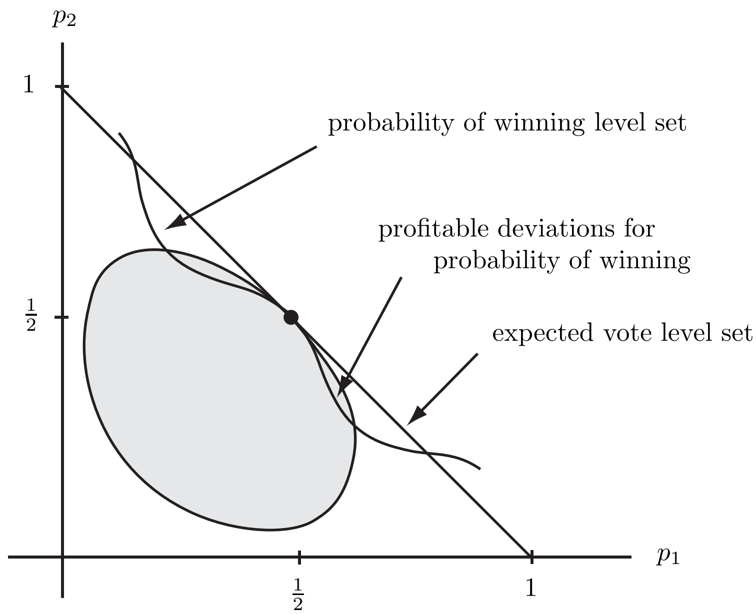

Figure 6: Difficulty with win motivation

of identically and uniformly distributed bias depicted in Figure 5, the fact that candidate A's expected vote level set is tangent to the set of achievable probability vectors implies that it is, indeed, optimal for the candidate to locate at the utilitarian point given that B does as well. In this case, the probability vector determined in equilibrium is $(\frac{1}{2}, \ldots, \frac{1}{2})$, and each candidate wins with probability one half. However, the set of probability vectors that deliver a probability of winning equal to one half does not have the same linear structure as we saw in expected vote maximization when there are three or more voters; Figure 6 provides a suggestive illustration of this, as the set of probability vectors delivering a probability of winning equal to one half "dips down" below the expected vote level set through $\overline{x}$.[14] This creates the possibility of deviations by candidate A (in fact, ones arbitrarily close to $\overline{x}$) that are profitable under win motivation, even though they reduce the candidate's expected vote. We will see that this is not a mere impression but that it is a real possibility in the probability of winning framework.

There is, nevertheless, a close connection between the equilibria generated by the two objective functions. Our first result, which extends results of Duggan (2000a) and Patty (2005), establishes that when biases are iid, 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

[14]The figure is "suggestive" only, because the depicted pathology does not occur with $n = 2$ voters.

candidates to take identical positions at a central point in the policy space. A consequence of Theorem 10 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.

Theorem 10 (Duggan; Patty) *In the stochastic partisanship model, assume win motivation; and assume biases θ_i are independently and identically distributed across voters. 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^*). Candidate A's maximization problem, given B's platform, is then

$$\max_{x \in X} \Pi_W(F_1(u_1(x) - u_1(x_B^*)), \ldots, F_n(u_n(x) - u_n(x_B^*)))$$

Since x_A^* is a best response for candidate A, it satisfies the necessary first order condition,

$$\sum_i \frac{\partial \Pi_W}{\partial p_i}(p^*) f_i(u_i(x_A^*) - u_i(x_B^*)) \nabla u_i(x_A^*) \quad = \quad 0, \tag{2}$$

where $p^* = (P_1(x_A^*, x_B^*), \ldots, P_n(x_A^*, x_B^*))$ is the vector of vote probabilities given the candidates' equilibrium platforms. Now consider the sum of voter utilities

$$\tilde{W}(x) \quad = \quad \sum_i \frac{\partial \Pi_W}{\partial p_i}(p^*) f_i(u_i(x_A^*) - u_i(x_B^*)) u_i(x),$$

with weights equal to the product of partial derivatives of the probability of winning function and the bias densities evaluated at the utility differences. Viewing these weights as fixed, $\tilde{W}$ is the sum of strictly concave functions of x (multiplied by strictly positive weights), and so $\tilde{W}$ is strictly concave. Furthermore, the first order condition for maximizing $\tilde{W}(x)$ is exactly the expression in (2), so x_A^* solves the first order condition for $\tilde{W}$, so it is the unique maximizer of $\tilde{W}$. A similar analysis for B leads to the same conclusion. In particular, $x_A^* = x_B^*$, which implies that the utility differences are zero. Using the assumption that biases are identically distributed, we can write $F_i(0) = \alpha$, and then the vector of vote probabilities is $p^* = (\alpha, \ldots, \alpha)$, and by symmetry of the probability of winning function, we have $\frac{\partial \Pi_W}{\partial p_i}(p^*) = \frac{\partial \Pi_W}{\partial p_j}(p^*)$ for all voters i and j. Since these partial derivatives are constant across voters, we can extract them from the summation in $\tilde{W}(x)$, and we conclude that $x_A^* = x_B^* = \overline{x}$, as required.

When is it an equilibrium for the candidates to locate at the utilitarian point? The next result gives partial answer to this question: a simple condition in terms of the second derivatives of voters' utility functions precludes profitable deviations near $\overline{x}$, ruling out the problem illustrated in Figure 6. Assuming, for simplicity, that biases are identically and 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, i.e., in one dimension, the second derivative of u_i at $\overline{x}$ is negative. Since we maintain the assumption that utilities are concave, the Hessian is already negative semi-definite, so the added restriction of negative definiteness appears unobjectionable. The result does not deliver the existence of a global equilibrium, but only a "local" one that 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. The statement here is adapted from Duggan (2000a), whereas Patty (2005) provides an extension to a more general model of voting and to multiple candidates.[15] Note that the result relies on the assumption that the bias density is positive at zero, which is innocuous, and that the derivative of the density is zero at zero; the latter is satisfied if the density is symmetric around zero, and it is satisfied if zero is the mode of the distribution.

Theorem 11 (Duggan; Patty) *In the stochastic partisanship model, assume win motivation; assume biases θ_i are independently distributed; assume voter utilities are twice continuously differentiable; and assume $f_i(0) > f_i'(0) = 0$ for all voters i. 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.*

To prove the result, I will borrow notation from the proof of Theorem 10. Of course, candidate A's maximization problem, given that B locates at the utilitarian point, is

$$\max_{x \in X} \Pi_W(F_1(u_1(x) - u_1(\overline{x})), \ldots, F_n(u_n(x) - u_n(\overline{x}))),$$

and the first order condition is as in (2) with $x_B^* = \overline{x}$. Note that $x_A = \overline{x}$ satisfies the first order condition, for then symmetry of the probability of winning function implies that the partial derivatives $\frac{\partial \Pi_W}{\partial p_i}$ are independent of i and can be dropped from the expression, leaving $\sum_i f_i(0) \nabla u_i(\overline{x}) = 0$, and this is just the first order condition for the utilitarian problem, which $\overline{x}$ satisfies by construction. Letting

$$\phi_i(x) \quad = \quad \frac{\partial \Pi_W}{\partial p_i}(p) f_i(u_i(x) - u_i(\overline{x})),$$

the first order condition for candidate A can be written as $\sum_i \phi_i(x) \nabla u_i(x) = 0$, where as we have noted, $\sum_i \phi_i(\overline{x}) \nabla u_i(\overline{x}) = 0$. The second order condition

[15] Whereas Patty (2005) assumes the Hessian of the expected vote objective is negative definite, I prove this from primitive assumptions. See also Aranson, Hinich, and Ordeshook (1974) for earlier results on the equivalence of vote and win motivation.

involves apparently unruly terms,

$$\begin{aligned}\frac{\partial \phi_i}{\partial x_k}(x) &= \sum_j \left[\frac{\partial^2 \Pi_W}{\partial p_i \partial p_j}(p) f_j(u_j(x) - u_j(\overline{x}))\frac{\partial u_j}{\partial x_k}(x)\right. \\ &\quad \left. + \frac{\partial \Pi_W}{\partial p_i}(p) f_i'(u_i(x) - u_i(\overline{x}))\frac{\partial u_i}{\partial x_k}(x)\right],\end{aligned}$$

that simplify greatly when evaluated at $x = \overline{x}$. Since we assume $f_i'(0) = 0$ for all voters, half of these terms drop out immediately. Furthermore, by symmetry of the probability of winning function, the cross partial derivatives of Π_W are positive and independent of i, so we may take them outside the summation. But the first order condition for the utilitarian point implies

$$\sum_j f_j(0)\frac{\partial u_j}{\partial x_k}(\overline{x}) = 0$$

for all $k = 1, \ldots, d$, and we conclude that $\frac{\partial \phi_i}{\partial x_k}(\overline{x}) = 0$. With the help of this observation, we can ignore all terms involving derivatives of ϕ_i, and the second derivative of candidate A's objective function evaluated at $\overline{x}$ reduces to

$$\sum_i \phi_i(\overline{x}) \begin{bmatrix} \frac{\partial^2 u_i}{\partial x_1^2}(\overline{x}) & \cdots & \frac{\partial^2 u_i}{\partial x_1 \partial x_d}(\overline{x}) \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 u_i}{\partial x_d \partial x_1}(\overline{x}) & \cdots & \frac{\partial^2 u_i}{\partial x_d^2}(\overline{x}), \end{bmatrix},$$

or in other words, it is the sum of the Hessian matrices of the voters evaluated at $\overline{x}$. Since each voter's utility function is concave, it follows that all of these matrices are negative semi-definite. And since we assume that the Hessian matrix of at least one voter i is negative definite, and since $\phi_i(\overline{x}) > 0$, it follows that the sum of matrices is negative definite. Thus, $\overline{x}$ satisfies the first and second order sufficient conditions for a strict local maximum, which implies that the candidate's probability of winning decreases at policies near $\overline{x}$, as required.

Interestingly, the added restriction of negative definiteness is needed for Theorem 11: the assumption of strict concavity alone is not enough for the result. Duggan (2000a) gives a three-voter example of the stochastic partisanship model with $X = [0, 1]$ and biases identically and uniformly distributed on $[-2.5, 2.5]$ in which the utilitarian point is not a local equilibrium:

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

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 11. 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 12 *In the stochastic partisanship model, assume win motivation. There exists 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.

4.3 Policy and Mixed Motivation

The stochastic partisanship model has not received much, if any, formal analysis under the assumption of policy motivation, but we can work out some basic properties of a simple, symmetric, one-dimensional model with one voter. Dropping voter subscripts, and letting $w \geq 0$ be the office benefit, candidate A's expected utility is

$$EU_A(x_A, x_B) = F(u(x_A) - u(x_B))(u_A(x_A) + w) + (1 - F(u(x_A) - u(x_B)))u_A(x_B).$$

We let $\tilde{x}_A$ and $\tilde{x}_B$ denote the ideal points of the candidates, and we assume that $\tilde{x}_A < \tilde{x}_B$. In contrast to the office motivated model, a candidate now faces a trade off between choosing a policy that is better in terms of policy utility (but less effective in winning the election) and choosing a policy that produces a higher probability of winning (but a worse policy outcome in that case). The analysis will focus on necessary first order conditions, but we will provide conditions that ensure sufficiency as well.

Before proceeding, we make the simple observation for the case of pure policy motivation, i.e., $w = 0$, that the candidates must locate at distinct platforms in equilibrium. This result does not depend on the assumption of a representative voter, and it contrasts with the results of Theorems 10 and 7 and Corollary 2 for office-motivated candidates.

Theorem 13 *In the stochastic partisanship, assume policy motivation. If the platform pair (x_A^*, x_B^*) is an equilibrium, then the candidates do not locate at the same policy position: $x_A^* < x_B^*$.*

To prove this result, simply note that whenever the two candidates' platforms are reversed, i.e., $x_A^* \geq x_B^*$, at least one candidate can deviate profitably by moving toward his or her ideal policy. Indeed, supposing $\tilde{x}_A < x = x_B^* \leq x_A^*$, let y lie strictly between $\tilde{x}_A$ and x. By concavity, y is strictly preferred to x by candidate A, and x is weakly preferred to x_A^* by the candidate. Our assumption that the support of F_i contains all possible utility differences implies that candidate A will win with positive probability after deviating to y. In particular, we have $u_A(y) > u_A(x)$ and $F_i(u_i(y) - u_i(x)) > 0$ for every voter. The latter implies $P(y, x) > 0$, and therefore

$$EU_A(y, x_B^*) \;=\; P(y,x)u_A(y) + (1 - P(y,x))u_A(x) \;>\; u_A(x) \;\geq\; EU_A(x_A^*, x_B^*),$$

so the deviation increases the candidate's expected utility, contradicting the assumption that (x_A^*, x_B^*) is an equilibrium. We conclude that $x_A^* < x_B^*$, as required.

To investigate the model with mixed motives, consider a one-dimensional model in which the policy space is the unit interval, $X = [0, 1]$, candidates A and B have ideal points $\tilde{x}_A = 0$ and $\tilde{x}_B = 1$, and candidates' utilities are inversely related, i.e., $u_A(x) = u_B(1 - x)$. Moreover, assume that the voter's utility function is symmetric around $1/2$, i.e., $u(x) = u(1 - x)$, and that the density of the bias is symmetric around zero, i.e., $f(x) = f(1-x)$. I refer to this as the *symmetric representative voter stochastic partisanship model.* Given that we already assume the candidates' policy utility functions are strictly concave,[16] the key condition is to impose some minimal structure on the distribution F of the voter's bias: we assume F is log concave, i.e., that $\ln(F(x))$ is concave. This condition, which simply means that $f(x)/F(x)$ is decreasing, is satisfied by many distributions of interest, including the normal and logistic distributions.

We will focus on equilibria (x_A^*, x_B^*) that are symmetric in the sense that $x_B^* = 1 - x_A^*$. In any interior equilibrium, candidate A's first order condition,

$$-F(u(x) - u(x_B^*))u_A'(x) \quad = \quad f(u(x) - u(x_B^*))u'(x)(u_A(x) + w - u_A(x_B^*))$$

is satisfied at $x = x_A^*$. Here, the lefthand side represents the marginal loss in utility incurred by candidate A from choosing a higher (and worse) policy platform; this effect only operates when the candidate wins, so it is weighted by the probability of winning. The righthand side represents the marginal effect on the probability of winning; this probability increases at the rate $f(u(x) - u(x_B^*))u'(x)$,

[16]The analysis to follow goes through under the more general assumption of log concavity, i.e., $\ln(u_A(x))$ and $\ln(u_B(x))$ are concave. This condition ensures that for all y, $u_A'(x)/(u_A(x) - u_A(y) + w)$ is decreasing on $[0, y]$, and similarly for B.

reflecting the density of the marginal voter (who is just indifferent, given his or her bias), and the return to higher probability of winning (instead of the other candidate) is the difference in the candidate's utility from the platforms plus the benefit of holding office. Note the immediate implication that the candidates cannot locate at their ideal points in equilibrium, for then $u'_A(0) = u'_B(1) = 0$, and the marginal loss from moving to a more moderate policy is zero, while the marginal benefit is positive.

The next theorem establishes existence and uniqueness of an equilibrium in the symmetric representative voter stochastic partisanship model.

Theorem 14 *In the symmetric representative voter stochastic partisanship model, assume mixed motivation; assume X is one-dimensional; and assume the distribution F of the voter's bias is log concave. There is a unique symmetric equilibrium (x^*_A, x^*_B); in this equilibrium, $0 < x^*_A \leq \frac{1}{2} \leq x^*_B < 1$, with strict inequalities if $w = 0$.*

To prove the result, note that $x^*_A \leq x^*_B$ holds, for otherwise either candidate could increase her payoff by deviating to the other's position, and we have argued that the candidates cannot locate at their ideal points. Strict inequalities hold for pure policy motivation by Theorem 13. For uniqueness, it is helpful to manipulate candidate A's first order condition, making the substitution $x^*_B = 1 - x$, as follows:

$$\frac{f(u(x) - u(1-x))u'(x)}{F(u(x) - u(1-x))} \quad = \quad -\frac{u'_A(x)}{u_A(x) + w - u_A(1-x)}. \tag{3}$$

The above equation is no longer the first order condition of a candidate's maximization problem, but it is satisfied at any symmetric equilibrium $(x, 1-x)$. By strict concavity, the righthand side of the equation is strictly increasing on X. Furthermore, by log concavity of F and concavity of u, the lefthand side is decreasing on $[0, \frac{1}{2}]$. Therefore, equality can hold for at most one value of x between zero and one half. Note that the lefthand side is positive at zero; and by $u'_A(0) = 0$, the righthand side equals zero at zero. The lefthand side equals zero at one half, while the righthand side approaches a positive number (if $w > 0$) or infinity (if $w = 0$) at zero. By concavity and differentiability, $u'_A(x)$ is continuous, so the both sides of the equality are continuous, and we conclude that there is exactly one solution to the equation on $[0, \frac{1}{2}]$, say x^*, and therefore $(x^*, 1-x^*)$ is the only possible symmetric equilibrium. To argue that x^* is, indeed, a best response to $1 - x^*$ for candidate A, we return to the first order condition for A with B fixed at $1 - x^*$:

$$\frac{f(u(x) - u(1-x^*))u'(x)}{F(u(x) - u(1-x^*))} \quad = \quad -\frac{u'_A(x)}{u_A(x) + w - u_A(1-x^*)}.$$

By construction, x^* satisfies the first order condition, and again the righthand side is strictly increasing (by concavity of u_A) and that the lefthand side is decreasing on $[0, \frac{1}{2}]$ (by concavity of u and log concavity of F), so there is no other solution to the candidate's first order condition. Therefore, x^* is a best response for A to $1 - x^*$, and by symmetry $1 - x^*$ is a best response for B to x^*, and therefore $(x^*, 1 - x^*)$ is an equilibrium, as required.

Uniqueness, if only in a simple example, is a useful property, as it facilitates the analysis of the comparative statics of the model. For example, in equation (3), which fully characterizes the equilibrium of the model, it is clear that an increase in w leads to a decrease in the righthand side, which necessitates an increase in x^* to maintain the equilibrium condition. In other words, when the benefit of holding office increases, the candidates become more willing to trade off worse policy outcomes for higher probability of winning, and they move toward the voter's ideal point. As office benefit becomes arbitrarily high, the equilibrium positions of the candidates converge toward the ideal point of the voter. This is intuitive: as w increases without bound, winning becomes arbitrarily more important than policy, and the objective functions of the candidates approximate the probability of winning objective; so it makes sense that their platforms should converge to the equilibrium under probability of winning maximization. For later comparison, note that even when office benefit is arbitrarily high, the candidates' positions do not actually reach the voter's ideal point, although they approach it.

The above analysis if of course incomplete. The focus has been only on symmetric equilibria and in a model that is greatly simplified for exposition. The obvious question is how the analysis changes when there are multiple voters, and the probability of winning function is less well-behaved; I leave this question open. Finally, I note that, as in the case of office-motivated candidates, the DFG theorem yields the general existence of a mixed strategy equilibrium.

Theorem 15 *In the stochastic partisanship model, assume policy or mixed motivation. There exists a mixed strategy equilibrium.*

Under mixed motivation, considering mixed strategy equilibria as the office benefit w becomes large, a simple continuity argument shows that if the mixed strategies of the candidates become close to degenerate (so they closely approximate pure strategies), then the limiting platform pair (x_A, x_B) is an equilibrium of under the probability of winning objective; and by Theorem 10, it follows that these platforms must coincide with the utilitarian point under general conditions.

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 policy preferences given by utility function $u_i(x, \theta_i)$, where θ_i is a preference parameter lying in a set Θ, which may either be a subset of the real line or possibly a subset of multidimensional space. We assume that for all θ_i, $u_i(\cdot, \theta_i)$ is strictly concave and differentiable. The parameter space Θ may be open or closed or, more generally, have an appropriate measurable structure; and we assume that for all x, $u_i(x, \cdot)$ is appropriately well-behaved over this space.[17] Given θ_i, the voter has a unique ideal point, denoted $\tilde{x}_i(\theta_i)$. 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 all distinct policies x and y, we have

$$\Pr(\{\theta \mid \text{for all } i,\, u_i(z, \theta_i) > u_i(y, \theta_i) > u_i(x, \theta_i)\}) \quad > \quad 0,$$

where z is the midpoint between x and y, and we have for all voters i,

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

and finally, for all voters i, we have

$$\Pr(\{\theta_i \mid \tilde{x}_i(\theta_i) = x\}) \quad = \quad 0.$$

We let H_i denote the distribution of voter i's ideal policy induced by G_i. 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

[17] Specifically, we assume that Θ is a Borel set and that $u_i(x, \cdot)$ is measurable with respect to the Borel structure on Θ.

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

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

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 and mixed motivation.

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 "average distribution" defined by $H_\alpha(x) = \frac{1}{n}\sum_i H_i(x)$. Our dispersion assumption implies that H_α is continuous and strictly increasing, so this average distribution has a unique median, denoted x_α. The next result, which is proved in Duggan (2006a), 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 16 (Duggan) *In the stochastic preference model, assume vote motivation; and assume X is unidimensional. 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 16, 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.[18] In the quadratic version of the model, x_γ is equivalent to a median in all directions,

[18] 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

in the usual sense.[19] 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.

Theorem 17 (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 than one half. 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.

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.

[19] 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 18 (Duggan and Jackson) *In the stochastic preference model, assume vote motivation. There exists a mixed strategy equilibrium.*

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 continuous and 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 19 (Calvert) *In the stochastic preference model, assume win motivation; and assume X is unidimensional. 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 16. Suppose toward a contradiction that there is an equilibrium in which some candidate locates at policy $x^* < x_\mu$. Since the game is constant sum, a standard interchangeability argument yields an equilibrium in which both candidates locate at x^*, in which case 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 17 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, existence of mixed strategy equilibria follows from results of Duggan and Jackson (2005).[20]

Theorem 20 (Duggan and Jackson) *In the stochastic preference model, assume probability of winning motivation. There exists a mixed strategy equilibrium.*

Moreover, Duggan and Jackson (2005) show that as voting behavior becomes close to Downsian, the mixed strategy equilibria must put probability arbitrarily close to one on policies close to the limiting median ideal policy. Thus, equilibria change in a continuous way when we add noise to the Downsian model.

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. In the one-dimensional model with quadratic voter utilities and in which the parameter θ_i is the ideal point of voter i, we know that the winner will be the candidate whose platform is closest to the median ideal point, θ_k. Thus, given platforms x_A and x_B with $x_A < x_B$, candidate A wins if $\frac{x_A+x_B}{2} < \theta_k$ or, after the toss of a coin, if $\theta_k < \frac{x_A+x_B}{2}$. By our dispersion assumption, the latter case occurs with probability zero, and we can ignore it; ties can only occur with positive probability if the candidates locate at the same position, $x_A = x_B$, in which case each wins with probability one half. Returning to the analysis when $x_A < x_B$, the probability candidate A wins is $H_\mu(\frac{x_A+x_B}{2})$, and we can write A's objective function as

$$EU_A(x_A, x_B) \quad = \quad H_\mu\left(\frac{x_A + x_B}{2}\right)u_A(x_A) + \left(1 - H_\mu\left(\frac{x_A + x_B}{2}\right)\right)u_A(x_B),$$

and likewise for candidate B.

The first result of this section, which does not rely on the assumption of quadratic utilities, was proved by Wittman (1983, 1990),[21] Hansson and Stuart (1984), Calvert (1985), and Roemer (1994). It establishes, consistent with our findings for the stochastic partisanship model, that the candidates can never locate at identical positions in equilibrium. Although the result was initially stated by Wittman and by Hansson and Stuart in the one-dimensional model, it is easily verified in the general multidimensional setting.

[20] This existence result is also established in Duggan (2007) by an alternative proof approach.

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

Theorem 21 (Wittman; Hansson and Stuart; Calvert; Roemer) *In the stochastic preference model, assume policy motivation; and assume X is unidimensional. If* (x_A^*, x_B^*) *is an equilibrium, then the candidates do not locate at the same policy position:* $x_A^* < x_B^*$.

The proof is very similar to that of Theorem 13. Suppose that $x_A^* \geq x_B^*$, and let one candidate deviate by moving toward his or her ideal policy. In particular, assuming without loss of generality that $\tilde{x}_A < x = x_B^* \leq x_A^*$, if candidate A moves to a policy y strictly between $\tilde{x}_A$ and x, then we have $u_A(y) > u_A(x) \geq u_A(x_A^*)$ and $\Pr(\{\theta_i \mid u_i(y, \theta_i) > u_i(x, \theta_i)\}) > 0$ for every voter, the latter implying $P(y, x) > 0$. Therefore

$$EU_A(y, x_B^*) \;=\; P(y,x)u_A(y) + (1 - P(y,x))u_A(x) \;>\; u_A(x) \;\geq\; EU_A(x_A^*, x_B^*),$$

so the deviation increases the candidate's expected utility, contradicting the assumption that (x_A^*, x_B^*) is an equilibrium, as required.

By Theorem 21, 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. We will return to this issue of equilibrium existence later.

Theorem 21 leaves open the question of whether an equilibrium exists in the first place. Roemer (1997) establishes that in the one-dimensional model, equilibria exist quite generally: it is sufficient that the distribution H_μ of the median ideal point is log concave.[22,23] Whereas Roemer (1997) assumes strictly concave policy utilities of the candidates, Duggan and Fey (2011) show that log concave utilities are sufficient. For simplicity, we assume strict concavity here, and we import the proof of Duggan and Fey (2011).

Theorem 22 (Roemer; Duggan and Fey) *In the stochastic preference model, assume policy motivation; assume X is one-dimensional; assume voter utilities are quadratic; and assume* H_μ *is differentiable and log concave on* X. *There exists an equilibrium; in every equilibrium,* $\tilde{x}_A < x_A^* < x_B^* < \tilde{x}_B$.

[22] Roemer (1994) assumes that each candidate always has a unique best response to the other, which directly imposes the conditions needed to apply Brouwer's theorem to deduce existence of an equilibrium. Later work derives this from assumptions on primitives.

[23] Assuming H_μ is differentiable with density h, log concavity implies that $h(x)/H_\mu(x)$ is decreasing. Note that h is automatically continuous from the left, and a discontinuity from the right would create a kink in the distribution H_μ. Thus, the density h is continuous.

The proof of the theorem is straightforward but makes use of Brouwer's fixed point theorem. For each pair (x_A, x_B) of policy platforms, we can consider the set of platforms that maximize candidate A's expected utility given x_B, and similarly for candidate B. Recalling that $\tilde{x}_A < \tilde{x}_B$, it is clear that best response platforms will always lie in the interval $[\tilde{x}_A, \tilde{x}_B]$ of ideal points. Now suppose these best responses are always unique. Then, by continuity of the candidates' objectives, the best response function for candidate A, which associates to each platform pair (x_A, x_B) the unique best response platform, say $r_A(x_A, x_B)$, will be continuous; likewise for B. Then Brouwer's fixed point theorem implies that the best response functions have a fixed point, i.e., there are platforms x_A^* and x_B^* such that $x_A^* = r_A(x_A^*, x_B^*)$ and $x_B^* = r_B(x_A^*, x_B^*)$. Therefore, (x_A^*, x_B^*) is an equilibrium, delivering the desired result. Of course, continuity of the objectives implies that each candidate will have at least one best response to the platform of the other. Here, we will verify the key condition that these best responses are always unique. To this end, consider any $x_B \in [\tilde{x}_A, \tilde{x}_B]$. Of course, if $x_B = \tilde{x}_A$, then A's unique best response is to locate at his or her ideal point. Otherwise, if $\tilde{x}_A < x_B$, then the argument in the proof of Theorem 21 shows that candidate A's best responses lie in the half-open interval $[\tilde{x}_A, x_B)$. Since candidate A's payoffs are differentiable on this interval, every best response must satisfy the first order condition: letting h denote the density of H_μ, this is

$$-H_\mu\left(\frac{x_A + x_B}{2}\right) u_A'(x_A) \quad = \quad h\left(\frac{x_A + x_B}{2}\right)\left(\frac{u_A(x_A) - u_A(x_B)}{2}\right)$$

on $(\tilde{x}_A, x_B)$, and the first order condition at the endpoint $\tilde{x}_A$ is the inequality

$$-H_\mu\left(\frac{\tilde{x}_A + x_B}{2}\right) u_A'(\tilde{x}_A) \quad \geq \quad h\left(\frac{\tilde{x}_A + x_B}{2}\right)\left(\frac{u_A(\tilde{x}_A) - u_A(x_B)}{2}\right).$$

The latter case is impossible: since $u_A'(\tilde{x}_A) = 0$, the inequality implies that $h(\frac{\tilde{x}_A + x_B}{2}) = 0$, but this contradicts log concavity.[24] In the former case, the exact first order condition is equivalent to

$$\frac{h\left(\frac{x_A + x_B}{2}\right)}{2H_\mu\left(\frac{x_A + x_B}{2}\right)} \quad = \quad \frac{-u_A'(x_A)}{u_A(x_A) - u_A(x_B)}.$$

As we vary x_A in the interval $[\tilde{x}_A, x_B]$, log concavity of H_μ implies that the lefthand side above is decreasing; moreover, concavity of u_A implies that the righthand side is strictly increasing. Thus, there is at most one solution to the first order condition, as required.

A key to the above existence result for policy motivated candidates — and one that distinguishes it from the case of office motivated candidates — is that the candidates' objectives are continuous. As before, each voter i's probability

[24] By log concavity, $h(x)/H_\mu(x)$ is decreasing, and $h(\frac{\tilde{x}_A + x_B}{2}) = 0$ thus implies $h(x) = 0$ for all $x \geq \frac{\tilde{x}_A + x_B}{2}$, contradicting the assumption that H_μ is strictly increasing.

$P_i(x_A, x_B)$ is continuous in the positions of the candidates whenever x_A and x_B are distinct, with discontinuities appearing only when one candidate crosses over another. Now, however, this form of discontinuity does not translate into a discontinuity in expected utilities: if we increase candidate A's policy, say x, toward B's, say y, then A's probability of winning increases to some number, say p, and A's expected utility satisfies

$$\begin{aligned}\lim_{x\to y} EU_A(x,y) &= \lim_{x\to y} P(x,y)u_A(x) + (1-P(x,y))u_A(y)\\ &= pu_A(y) + (1-p)u_A(y)\\ &= u_A(y)\\ &= EU_A(y,y).\end{aligned}$$

Thus, even though p may not equal one half, A's expected utility varies continuously with x_A, so 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 (2005) show that if the distribution of voting parameters piles probability mass near a particular vector $(\theta_1, \ldots, \theta_n)$, and if there is no equilibrium in the Downsian model (with voter preferences given by the vector $(\theta_1, \ldots, \theta_n)$) in the game between policy motivated candidates, then there does not exist an equilibrium in the probabilistic voting model. From Theorem 5, we know that equilibria in the Downsian model with policy motivated candidates will sometimes fail to exist; and when the policy space has dimension three or more, this existence failure will actually be typical. We conclude that in higher dimensions, equilibria in the stochastic preference model can fail to exist; and this will be the typical situation when the distribution of parameters is sufficiently concentrated, i.e., the candidates have a large amount of information about voter preferences.

What is more surprising is that non-existence of equilibria can occur even in the one-dimensional model, if the assumption of log concavity of H in Theorem 22 is relaxed. Duggan and Fey (2011) demonstrate existence failure in a highly structured setting — they assume quadratic utilities for the candidates and a symmetric, single-peaked density for the distribution of ideal policies.[25] They assume $X = [0,1]$, quadratic utilities with $\tilde{x}_A = 0$ and $\tilde{x}_B = 1$, and a piece-wise linear distribution G that puts probability arbitrarily close to one near 1/2. The problem is that the candidates' best responses are not uniquely defined, creating the possibility of a "jump" in their reaction functions. In Figure 7, 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

[25] 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.

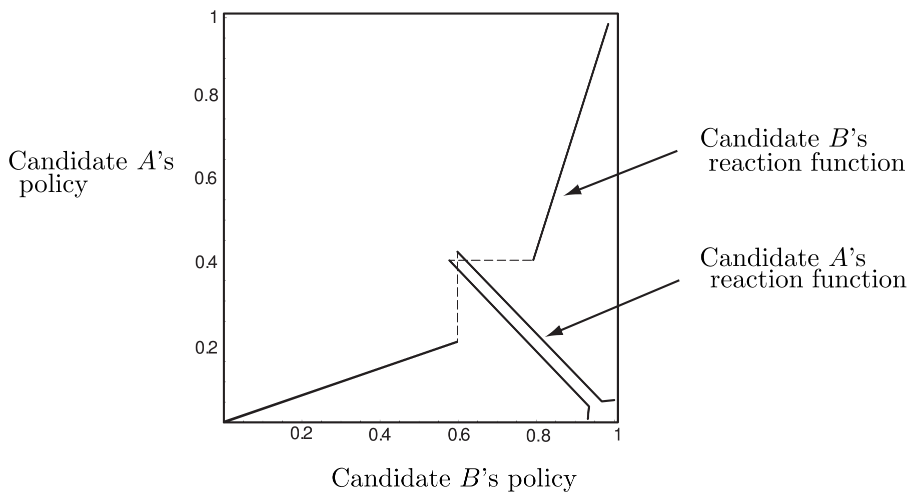

Figure 7: Equilibrium non-existence with policy motivation

(slightly greater than .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.[26] As we have discussed, the candidates' objective functions are continuous in the model with policy motivation, so the existence failure here is solely due to the violation of quasi-concavity of the candidates' expected utilities.

As a consequence, 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 23 *In the stochastic preference model, assume policy motivation. There exists a mixed strategy equilibrium.*

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

[26] 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.

behavior of voters in the Downsian model, we have seen that pure strategy equilibria can fail to exist. By Theorem 23, 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 limiting median ideal policy. Thus, equilibria change in a continuous way when we add noise to the Downsian model.

Theorem 22 has no implications for uniqueness, nor for existence in the model with mixed motives, where candidates attach a positive benefit $w > 0$ to holding office. The latter topic is covered in the next subsection, and general results for uniqueness are not available. Adding more structure to the model, however, we can establish uniqueness among the class of symmetric equilibria. We now consider the *symmetric stochastic partisanship model*, in which X is the interval $[-1, 1]$, the candidates' utilities are symmetric around zero, i.e., $\tilde{x}_A < 0 < \tilde{x}_B$ with $u_A(x) = u_B(-x)$, and the distribution of the median voter is also symmetric around zero, i.e., $H_\mu(x) = 1 - H_\mu(-x)$. Note that former assumption implies that the ideal points of the candidates are symmetric, so $\tilde{x}_A = -\tilde{x}_B = \tilde{x}$, and the latter assumption implies symmetry of the density, i.e., $h(x) = h(-x)$. The next result is a special case of a result of Bernhardt, Duggan, and Squintani (2009), who consider the more general case of mixed motivation,[27] which is presented in the next subsection; we defer the proof until then.

Theorem 24 (Bernhardt, Duggan, and Squintani) *In the symmetric stochastic preference model, assume policy motivation; assume X is unidimensional; assume voter utilities are quadratic; and assume H_μ is differentiable and log concave on X. There is a symmetric equilibrium; it is unique among the class of symmetric equilibria, and it is the platform pair $(x^*, -x^*)$, where x^* is the unique negative solution to*

$$-\frac{u'_A(x)}{u_A(x) - u_A(-x)} \quad = \quad h(0).$$

Thus, with appropriate symmetry assumptions, the conditions of Theorem 22 are sufficient for uniqueness among the class of symmetric equilibria. When candidate utilities are quadratic, the equilibrium platform x^* of candidate A can be solved for explicitly as

$$x^* \quad = \quad \frac{\tilde{x}}{1 - 2h(0)\tilde{x}}.$$

As mentioned, log concavity of H_μ is quite general and is satisfied, for example, by the normal distribution. For the normal special case, we have

$$x^* \quad = \quad \frac{\tilde{x}\sigma}{\sigma - \tilde{x}\sqrt{2/\pi}},$$

[27] As well, Bernhardt, Duggan, and Squintani (2009) use a condition slightly weaker than log concavity of H_μ at the cost of a more complicated proof.

where σ is the variance of the distribution of the median ideal point. And in the uniform special case, if we assume H_μ is uniform on the interval $[\tilde{x}, -\tilde{x}]$ of the candidates' ideal points, then we have simply $x^* = \frac{\tilde{x}}{2}$. Consistent with Theorem 13 for the stochastic partisanship model, but in contrast to Theorem 4 for the Downsian model, we see that the candidates' platforms separate in equilibrium. In contrast with Theorem 13, we have a new comparative static with respect to the density of H_μ at zero: as the distribution of voter preferences becomes more concentrated (so $h(0)$ increases, or in the normal model σ decreases), we see that x^* increases toward zero; in effect, as the likelihood that the median voter is located near zero increases, the candidates' become more willing to tradeoff policy for probability of winning, pulling them toward the center of the policy space. Similarly, as the candidates become more centrist (so $\tilde{x}$ increases toward zero), equilibrium platforms become more moderate, reflected in an increase in x^*.

5.4 Mixed Motivation

We now add a positive benefit $w > 0$ of holding office to obtain the model with mixed motives:

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

and likewise for candidate B. Although apparently innocuous, the addition of these office term creates technical difficulties. First, the objective functions of the candidates are again discontinuous: as we vary candidate A's policy x_A, keeping x_B fixed, A's payoff jumps discontinuously as x_A crosses over x_B (unless x_B equals the median x_μ of the distribution H_μ). Second, and related, we can no longer be sure that candidate A's optimal policy choice is not to the far side of B's platform; specifically, assuming $\tilde{x}_A < x_\mu < \tilde{x}_B$, if candidate B locates close to A's ideal point, it may be beneficial for A to locate between B's platform and the median of H_μ in order to obtain the benefits of office with higher probability. This is not merely a technical glitch: Ball (1999) gives a one-dimensional example with uniform distribution H_μ in which an equilibrium fails to exist due to these problems.

The proof of Theorem 22 does, however, yield a platform pair (x_A^*, x_B^*) such that x_A^* maximizes candidate A's expected utility over the policies to the same side of x_B^* as her ideal point. In particular, the platform established there is a local equilibrium.

Theorem 25 (Duggan and Fey) *In the stochastic preference model, assume mixed motivation; assume X is one-dimensional; and assume H_μ is log concave. There exists a local equilibrium.*

In the symmetric model, these problems, while still potentially a concern, do not interfere with the existence and uniqueness of symmetric equilibrium; and in fact, the characterization of Theorem 24 carries over straightforwardly. The one new feature introduced by mixed motives is the possibility that the candidates both locate at the median x_μ of the distribution of the median voter: if $|u'_A(0)| \leq wh(0)$, so that the marginal incentive of moving toward A's ideal point is outweighed by the marginal incentive of capturing the office benefit, then the candidates will locate at the median of medians, the equilibrium from Theorem 19 with probability of win motivation. Clearly, this contrasts with the stochastic partisanship result in Theorem 13, in which no amount of office benefit induces the candidates to locate at the same policy.

Theorem 26 (Bernhardt, Duggan, and Squintani) *In the symmetric stochastic preference model, assume mixed motivation; assume X is unidimensional; assume voter utilities are quadratic; and assume H_μ is differentiable and log concave on X. There is a symmetric equilibrium; it is unique among the class of symmetric equilibria, and it is the platform pair $(x^*, -x^*)$, where x^* is defined as follows: if $|u'_A(0)| \leq wh(0)$, then $x^* = 0$; and if $|u'_A(0)| > wh(0)$, then x^* is the unique negative solution to*

$$-\frac{u'_A(x)}{u_A(x) + w - u_A(-x)} \quad = \quad h(0).$$

To prove the result, recall from the proof of Theorem 22 (after including the office benefit term) that given platform x_B for candidate B, the first order condition for candidate A is

$$\frac{h\left(\frac{x+x_B}{2}\right)}{2H_\mu\left(\frac{x+x_B}{2}\right)} \quad = \quad \frac{-u'_A(x)}{u_A(x) + w - u_A(x_B)}. \tag{4}$$

Given a symmetric equilibrium $(x^*, -x^*)$ with $x^* < 0$, we rule out $x^* = \tilde{x}_A$ by the argument in the proof of Theorem 22; thus, the first order condition is satisfied at $x = x^*$ with equality, and it follows that

$$h\left(0\right) \quad = \quad \frac{-u'_A(x^*)}{u_A(x^*) + w - u_A(-x^*)} \tag{5}$$

must hold. By strict concavity of u_A, the righthand side is strictly increasing, and we conclude that there is at most one symmetric equilibrium $(x^*, -x^*)$ with $x^* < 0$. Consider the case $|u'_A(0)| > wh(0)$. Then the righthand side of (5) equals zero when x^* equals A's ideal point $\tilde{x}_A$, and it approaches a positive number (if $w > 0$) greater than $h(0)$ or infinity (if $w = 0$) when x^* is close to zero. In this case, therefore, there is a unique solution to equation (5), and by the arguments in the proof of Theorem 22, the policy platform x^* is a best response for candidate A from the set of policies $x \leq -x^*$. The possibility that candidate A may gain from locating to the other side of B's position does not arise, because $x^* < x_\mu < -x^*$, and then policies greater than $-x^*$ generate lower policy utility and a lower probability of winning for A than policies less than (and close to) x^*, e.g., the policy $\frac{x^*}{2}$. Thus, the platform pair $(x^*, -x^*)$ is an equilibrium and is unique among the symmetric equilibria with $x^* < 0$. On the other hand, the pair $(0,0)$ is not an equilibrium, because the derivative of A's expected utility evaluated at zero, given $x_B = 0$, is

$$\frac{d}{dx_A} EU_A(0,0) \quad = \quad \frac{h(0)w - u'_A(0)}{2} \quad < \quad 0,$$

so candidate A could obtain a higher level of expected utility by decreasing his or her platform. We conclude that $(x^*, -x^*)$ is the unique symmetric equilibrium in this case. Finally, consider the case $|u'_A(0)| \leq wh(0)$. Now the righthand side of (5) is greater than or equal to $h(0)$ when $x^* = 0$, and since it is strictly increasing, there is no strictly negative solution, i.e., there is no symmetric equilibrium with $x^* < 0$. Moreover, $(0,0)$ is now an equilibrium. Indeed, candidate A cannot deviate profitably by moving to a platform greater than zero; and to see that the candidate cannot deviate profitably to a platform less than zero, note that any best response for A must satisfy the first order condition (4) with $x_B = 0$. But the lefthand side is greater than or equal to the righthand side when $x = 0$; and since the lefthand side is weakly decreasing, and the righthand side is strictly increasing, the first order condition does not hold at any platform $x_A < 0$. We conclude that $(0,0)$ is an equilibrium and, moreover, is unique among symmetric equilibria.

It is straightforward to solve explicitly for the unique symmetric equilibrium when more structure is imposed on candidate utilities or the distribution of the median voter's ideal point. Consistent with the analysis of Theorem 13, we see that as we increase office benefit from $w = 0$, the symmetric equilibrium policies from Theorem 24 under policy motivation move toward the center of the policy space, eventually coinciding at the median x_μ of the distribution H_μ. When the attractiveness of winning increases, the candidates trade off policy outcomes for probability of winning, pulling them toward more moderate policy platforms.

I end with a last remark on mixed strategy equilibrium. Despite the discontinuities noted by Ball (1999) and concomitant non-convexities, Ball (1999) proves existence of mixed strategy equilibrium in the one-dimensional model; the result of Duggan and Jackson (2005) holds in general policy spaces.

Theorem 27 (Ball; Duggan and Jackson) *In the stochastic preference model, assume mixed motivation. There exists a mixed strategy equilibrium.*

Similar to office and policy motivation, Duggan and Jackson (2005) show that 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 offered 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 (2008), 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.

Finally, I will make the (perhaps obvious) point that many topics of interest are not covered in this survey. I do not cover electoral models with a continuum of voters, which admittedly make up a significant portion of the electoral modeling literature. I also do not touch on models of valence (Groseclose (2001) and Schofield (2003)) where one candidate is perceived by all voters as superior to the other in some way, irrespective of policy positions; models of privately informed candidates (Bernhardt, Duggan, and Squintani (2007)) in which candidates receive private signals about the location of the voters' ideal points; models with multiple candidates; models of abstention, etc.

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 (2008) "A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces," *Quarterly Journal of Political Science,* 3: 269–299.

[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] D. Bernhardt, J. Duggan, and F. Squintani (2007) "Electoral Competition with Privately Informed Candidates," *Games and Economic Behavior,* 58: 1–29.

[11] D. Bernhardt, J. Duggan, and F. Squintani (2009) "The Case for Responsible Parties," *American Political Science Review,* 103: 570–587.

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

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

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

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

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

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

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

[19] 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.

[20] 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.

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

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

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

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

[25] J. Duggan (2006a) "Candidate Objectives and Electoral Equilibrium," in *The Oxford Handbook of Political Economy,* eds., B. Weingast and D. Wittman, Oxford: Oxford University Press.

[26] J. Duggan (2006b) "A Note on Uniqueness of Electoral Equilibrium When the Median Voter is Unobserved," *Economics Letters,* 92: 240–244.

[27] J. Duggan (2007) "Equilibrium Existence for Zero-sum Games and Spatial Models of Elections," *Games and Economic Behavior,* 60: 52–74.

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

[29] J. Duggan and M. Fey (2011) "Indeterminacy in Elections with Policy-Motivated Candidates and a Single Issue Dimension," mimeo.

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

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

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

[33] 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.

[34] T. Groseclose (2001) "A Model of Candidate Location when One Candidate Has Valence Advantage," *American Journal of Political Science,* 45: 862–886.

[35] 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.

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

[37] 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.

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

[39] 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.

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

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

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

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

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

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

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

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

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

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

[50] 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.

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

[52] J. Patty (2005) "Local Equilibrium Equivalence in Probabilistic Voting Models," *Games and Economic Behavior,* 51: 523–536.

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

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

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

[56] J. Roemer (1997) "Political-economic Equilibrium when Parties Represent Constituents: The Unidimensional Case," *Social Choice and Welfare,* 14: 479–502.

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

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

[59] A. Saporiti (2008) "Existence and Uniqueness of Nash Equilibrium in Electoral Competition Games: The Hybrid Case," *Journal of Public Economic Theory,* 10: 827–857.

[60] N. Schofield (2003) "Valence Competition in the Spatial Stochastic Model," *Journal of Theoretical Politics,* 15: 371–383.

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

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

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

[64] 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.