Equilibrium Existence in Discontinuous Zero-Sum Games with Applications to Spatial Models of Elections*

John Duggan
Department of Political Science
and Department of Economics
University of Rochester
Rochester, NY 14627

June 7, 2005

*I thank Jeff Banks, Matt Jackson, Michel Le Breton, Phil Reny, and Jeroen Swinkels for their input. Support from the National Science Foundation, grant number SES-0213738, is also gratefully acknowledged.

Abstract

A theorem on existence of mixed strategy equilibria in discontinuous zero-sum games is proved and is applied to three models of elections. First, the existence theorem yields a mixed strategy equilibrium in the multidimensional spatial model of elections with three voters. A nine-voter example shows that a key condition of the existence theorem is violated for general finite numbers of voters and illustrates an obstacle to a general result. Second, the theorem provides a simple and self-contained proof of Kramer's (1978) existence result for the multidimensional model with a continuum of voters. Third, existence follows for a class of multi-dimensional probabilistic voting models with discontinuous probability-of-winning functions.

1 Introduction

This paper contributes to a larger literature on existence of equilibrium in games with discontinuous payoff functions, including Dasgupta and Maskin (1986), Simon (1987), and Reny (1999), but it focuses on a special class of games, namely, two-player, zero-sum games. In that sense, it is related to papers by Sion (1958) and Mertens (1986), who weaken continuity of payoffs to upper semicontinuity in own strategies. The paper is organized around the following theorem. First, say a mixed strategy for a player is "hospitable" if it makes the other player's payoffs are continuous in his/her own strategy. Assume compact strategy sets. Assume that given any mixed strategy for one player that allows the other to achieve a positive payoff, the other player can achieve a positive payoff with a hospitable strategy. Then there is a mixed strategy equilibrium.

The paper provides a simple proof of this result, one that is based on an infinite-dimensional version of the separating hyperplane theorem and has a clear geometric interpretation. Alternatively, it is shown that the condition implies Reny's (1999) condition of better reply security, so the existence theorem follows from results of his paper. Beyond the proof approach, which may be of independent interest, the central contribution of this paper can then be viewed as the definition of several useful sufficient conditions for better reply security in symmetric, two-player, zero-sum games and the application of those sufficient conditions to three classes of electoral models, where discontinuities arise quite naturally. Thus, standard existence results for continuous games, such as Glicksberg (1952), cannot be applied. Furthermore, the upper semicontinuity conditions of Sion (1958) and Mertens (1986) are violated, and the conditions of the other above-cited papers on existence in discontinuous games have proved difficult to verify.

In the classical spatial model of elections, as developed by Black (1958), Downs (1957), and Hotelling (1929), two candidates take positions in a convex policy space before an election. A finite number of voters, who are not modelled as players, have continuous, strictly convex preferences and are assumed to vote for their preferred candidates.[1] The candidates are assumed

[1]Duggan and Jackson (2005) take an alternative approach to existence by incorporating voters as players, who, as such, may vote arbitrarily when indifferent between the candidates.

to have full information about voter preferences and to be office-motivated, seeking only to win the election rather than, say, to implement preferred policies. It is known that when the policy space is unidimensional (and voter preferences are therefore single-peaked), the electoral game has pure strategy equilibria. Moreover, a pair of positions is an equilibrium if and only if each candidate locates at a median with respect to the voters' ideal points. If the number of voters is odd, then there is a unique median, and the game actually has a unique equilibrium. When the policy space is multidimensional, a pair of positions is a pure strategy equilibrium if and only if each candidate locates at a position that is undominated according to majority rule, the set of such points being known as the "core" of the spatial voting game. A long line of work in spatial modelling has shown that for multidimensional policy spaces, nonemptiness of the core implies very restrictive conditions on voter preferences.[2] These necessary conditions are generically (in variously defined senses) violated, with the conclusion that pure strategy equilibria of the electoral game almost never exist.

It is true, of course, that mixed strategy equilibria exist if the policy space is finite. Indeed, Laffond, Laslier, and Le Breton (1993) show that in the absence of majority-ties, there is exactly one mixed strategy equilibrium in this setting. When the policy space consists of divisions of a "dollar" among the set of voters, and each voter simply prefers more of the dollar, the electoral game reduces to the Colonel Blotto game, and it is known that mixed strategy equilibria exist in a version of this game.[3] Banks, Duggan, and Le Breton (2002) have shown generally that if there is a mixed strategy equilibrium of the electoral game, then the equilibrium strategies of the candidates must have support contained in the uncovered set, a centrally located region of the policy space — but existence in the spatial model has remained elusive.

The first application of the existence theorem establishes existence of a mixed strategy equilibrium in the three-voter case of the classical spatial model. A nine-voter example shows that the verification of the key condition of the existence theorem becomes highly problematic: it is violated despite, in one version of the example, the presence of a mixed strategy equilibrium. An implication is that the condition is not necessary for existence. Reny's (1999) diagonal better reply security is satisfied, but to secure

[2]See Plott (1967), Schofield (1983), Banks (1995), and Saari (1997).

[3]See Gross and Wagner (1950), Borel (1953), and Laslier and Picard (2002).

deviating payoffs as called for in Reny's condition, pure strategies are not sufficient. Rather, mixed strategies are sometimes needed, and those mixed strategies must sometimes be "far" from the initial deviation being secured. We leave open the question as to whether Reny's result can be used to prove a general existence theorem for the Downsin model, but it is clear from the above-mentioned example that the verification of his condition would be challenging.

The second application is to a version of the spatial model in which there is a continuum of voters and candidates are "plurality-motivated," meaning that each maximizes his/her votes net of the other candidate's votes. It is known that in the unidimensional model, the median voter result extends to provide a pure strategy equilibrium. In multiple dimensions, these alternative assumptions rule out certain especially problematic discontinuities in the candidates' payoffs. Kramer (1978) uses this structure to prove existence of a mixed strategy equilibrium, but his proof is quite involved, and is not self-contained. It is shown below, however, that the assumptions of this paper's existence theorem are easily verified in the continuum model, yielding a straightforward path to existence. Moreover, we are able to shed much of the usual structure imposed on voter preferences, dropping transitivity conditions entirely and assuming only the weak convexity property of "star-shapedness." Though such generality is not often a high priority in voting models, it facilitates the derivation of our next result.

The third application is to spatial models in which candidates have incomplete information about the preferences of voters, so that voting behavior is a random variable from the perspective of the candidates. In the unidimensional model, a version of the median voter theorem is well-known: a pair of positions is a pure strategy equilibrium if and only if each candidate locates at a median of the distribution of the median voter's ideal point. In multiple dimensions, Calvert (1985) and others have noted extremely restrictive conditions on the distribution of voter preferences under which pure strategy equilibria exist. Existence of mixed strategy equilibrium in such general settings has not been considered,[4] but the theorem of this paper yields existence of a mixed strategy equilibrium. Interestingly, a separate proof is not

[4]An exception is Ball (1999), who proves existence in a specific version of the probabilistic voting model. He assumes a single dimension, uniformly distributed median, and candidates with "mixed" motives. Banks and Duggan (2005) examine mixed strategy equilibria in a class of probabilistic voting models with continuous candidate payoffs.

required: the argument given here exploits a formal equivalence between the probabilistic voting model and the general continuum model, as described above.

2 Existence in Zero-Sum Games

A *two-player, zero-sum game* is a tuple (X, X', μ, μ', u). Here, X and X' are nonempty subsets of Hausdorff topological spaces, μ and μ' are finite Borel measures on X and X' with $\mu(X) > 0$ and $\mu'(X') > 0$, the product $X \times X'$ is endowed with the product topology, and $u\colon X \times X' \to \Re$ is a bounded, Borel measurable payoff function. We say the game is *symmetric* if $X = X'$ and, for all $x, x' \in X$, we have $u(x, x') = -u(x', x)$. A pair (x, x') is a *pure strategy equilibrium* if for all $y \in X$ and all $y' \in X'$, we have $u(y, x') \leq u(x, x') \leq u(x, y')$.

A mixed strategy, denoted σ or σ', is a Borel probability measure on X or X'. Let Σ denote the set of all mixed strategies on X and Σ' the set of mixed strategies on X'. We extend the payoff function u to mixed strategy pairs as follows:

$$U(\sigma, \sigma') \;=\; \int\int u(x, x')\, \sigma(dx)\sigma'(dx').$$

If σ is degenerate on x, we write simply $U(x, \sigma')$ for this payoff. Because u is bounded, we may view this as defining a mapping $U(\cdot, \sigma')\colon X \to \Re$, which, by Fubini's theorem, is Borel measurable.[5] A pair (σ, σ') is a *mixed strategy equilibrium* if for all $y \in X$ and all $y' \in X'$,

$$U(y, \sigma') \;\leq\; U(\sigma, \sigma') \;\leq\; U(\sigma, y').$$

That existence of a mixed strategy equilibrium does not follow from the previous assumptions, even when strategy sets are assumed to be compact, is illustrated in the next section.

Define the subset of mixed strategies

$$\Sigma^c \;=\; \{\sigma' \in \Sigma' \mid U(\cdot, \sigma')\colon X \to \Re \text{ is continuous } \}$$

[5]When defining notation and terminology from the perspective of one player, we assume similar conventions for the other.

as the set of mixed strategies for the first player that make the second player's payoff continuous in his/her own strategy. We refer to such strategies as *hospitable.* Two conditions will play important roles in what follows.

(C0) The sets X and X' are compact.

(C1) For all $\sigma' \in \Sigma'$ and for all $x \in X$,

$$U(x, \sigma') > 0 \quad \Rightarrow \quad \sup_{\sigma \in \Sigma^c} U(\sigma, \sigma') > 0.$$

The first of these conditions is standard. The second condition stipulates that whenever the first player can achieve a positive payoff using a pure strategy against some mixed strategy of the other player, he/she can achieve a positive payoff using a hospitable strategy.

We will also make use of several measure-theoretic conditions that will imply (C1). We say σ is *absolutely continuous* if it is absolutely continuous with respect to μ, and we let Σ° denote the set of absolutely continuous mixed strategies on X.

(C2) For all $\sigma' \in \Sigma^{\circ\prime}$, the function $U(\cdot, \sigma') \colon X \to \Re$ is continuous, and for all $\sigma \in \Sigma^\circ$, the function $U(\sigma, \cdot) \colon X' \to \Re$ is continuous.

(C3) For all $\sigma' \in \Sigma'$ and for all $x \in X$,

$$U(x, \sigma') > 0 \quad \Rightarrow \quad \sup_{\sigma \in \Sigma^\circ} U(\sigma, \sigma') > 0.$$

(C4) For all $\sigma' \in \Sigma'$ and for all $x \in X$,

$$\sup_{\sigma \in \Sigma^\circ} U(\sigma, \sigma') \quad \geq \quad U(x, \sigma'),$$

and for all $\sigma \in \Sigma$ and for all $x' \in X'$,

$$\inf_{\sigma' \in \Sigma^{\circ\prime}} U(\sigma, \sigma') \quad \leq \quad U(\sigma, x').$$

Condition (C2) is a minimal continuity assumption requiring that when one candidate uses an absolutely continuous strategy, the other's payoff is continuous in his/her own strategy. Condition (C3) demands that whenever

the first player can achieve a positive payoff using a pure strategy against some mixed strategy of the other player, he/she can achieve a positive payoff using an absolutely continuous strategy. Condition (C4) strengthens this, saying that both players can approximate the payoff of any pure strategy against any mixed strategy of the opponent by using absolutely continuous strategies. The next proposition is immediate.

Proposition 1 *If (C2) and (C3) are satisfied, then (C1) is satisfied.*

The next proposition gives a simple sufficient condition for (C2) when strategy sets are subsets of Euclidean space. An implication is that (C2) is satisfied if for every strategy of a player, the player's payoff as a function of the opponent's strategy is continuous at all but a countable number of points.

Proposition 2 *Assume X and X' are subsets of finite-dimensional Euclidean space, and assume $\mu = \lambda$ and $\mu' = \lambda'$, where λ and λ' are Lebesgue measure on X and X', respectively. If for all $x \in X$ and $x' \in X$,*

$$\begin{aligned} \lambda'\left(\{z \in X' \mid u(\cdot, z) \text{ is discontinuous at } x\}\right) &= 0 \\ \lambda\left(\{z \in X \mid u(z, \cdot) \text{ is discontinuous at } x'\}\right) &= 0, \end{aligned}$$

then (C2) is satisfied.

For the proof of the proposition, fix $\sigma' \in \Sigma^\circ$, and consider a sequence $\{x_n\}$ converging to x. To see that the sequence of functions $\{u(x^n, \cdot)\}$ converges λ'-almost everywhere to $u(x, \cdot)$, note that

$$\lambda'\left(\{z \in X' \mid u(\cdot, z) \text{ is discontinuous at } x\}\right) = 0,$$

so, by absolute continuity,

$$\sigma'\left(\{z \in X' \mid u(\cdot, z) \text{ is discontinuous at } x\}\right) = 0.$$

Taking any $z \in X'$ such that $u(\cdot, z)$ is continuous at x, we have $u(x^n, z) \to u(x, z)$. Then $U(x^n, \sigma') \to U(x, \sigma')$ follows from Lebesgue's dominated convergence theorem. We conclude that (C2) is satisfied.

The next proposition gives a simple sufficient condition for (C4), and therefore for (C3), when X and X' are subsets of Euclidean space.

Proposition 3 *Assume X and X' are subsets of finite-dimensional Euclidean space, and assume $\mu = \lambda$ and $\mu' = \lambda'$, where λ and λ' are Lebesgue measure on X and X', respectively. Suppose that*

- *for all $\sigma' \in \Sigma'$ and all $x \in X$ such that $U(x, \sigma') > 0$, there exists a nonempty open set $G \subseteq X$ such that for all $z \in G$, $U(z, \sigma') > 0$, and*
- *for all $\sigma \in \Sigma$ and all $x' \in X'$ such that $U(\sigma, x') < 0$, there exists a nonempty open set $G \subseteq X'$ such that for all $z \in G$, $U(\sigma, z) < 0$.*

Then (C4) is satisfied.

The proof is immediate: fix $\sigma' \in \Sigma'$ and $x \in X$ such that $U(x, \sigma') > 0$, and let σ be the uniform distribution on any nonempty, open ball contained in the open set G given in the proposition.

The applications considered in the remainder of this paper are organized around the following theorem on existence of mixed strategy equilibria in symmetric, two-player, zero-sum games.

Theorem 1 *In any symmetric, two-player, zero-sum game satisfying (C0) and (C1), there exists a mixed strategy equilibrium.*

For the proof of the theorem, let

$$\mathcal{U} \;=\; \{U(\cdot, \sigma') \mid \sigma' \in \Sigma^c\}$$

denote the set of payoff functions for one player induced by hospitable strategies of the other player. Note that this set is non-empty: if u is constant at zero, then $\Sigma^c = \Sigma$; otherwise, there exist $x, x' \in X = X'$ such that $u(x, x') > 0$, and non-emptiness follows from (C1). Further, $\mathcal{U}$ is convex, since convex combinations of hospitable strategies are hospitable. By definition, $\mathcal{U}$ is a subset of $C(X)$, the vector space of real-valued continuous functions on X. Let

$$\mathcal{N} \;=\; \{f \in C(X) \mid \text{ for all } x \in X, f(x) \leq 0\}$$

denote the non-positive orthant of $C(X)$, which is also convex. Note that if $U(\cdot, \hat{\sigma}) \in \mathcal{U} \cap \mathcal{N}$, then $(\hat{\sigma}, \hat{\sigma})$ is a mixed strategy equilibrium: $U(\hat{\sigma}, \hat{\sigma}) = 0$ by symmetry, and then

$$U(y, \hat{\sigma}) \;\leq\; 0 \;\leq\; U(\hat{\sigma}, y)$$

for all $y \in X$. Suppose, then, that $\mathcal{U} \cap \mathcal{N} = \emptyset$. The function that is identically negative one is an internal point of $\mathcal{N}$, and the separating hyperplane theorem (Aliprantis and Border, 1999, Theorem 5.46) yields a non-zero linear functional $p\colon C(X) \to \Re$ that properly separates $\mathcal{U}$ and $\mathcal{N}$, i.e., for all $f \in \mathcal{N}$ and all $f' \in \mathcal{U}$, $p(f) \leq p(f')$, with strict inequality for at least one f and one f'. Since the function that is identically zero is in $\mathcal{N}$, we have $0 \leq p(f')$ for all $f' \in \mathcal{U}$, and it follows that $p(f) \leq 0$ for all $f \in \mathcal{N}$. Therefore, p is a positive linear function. By (C0), each $f' \in C(X)$ has compact support, so the Riesz representation theorem (Rudin, 1987, Theorem 2.14) yields a Borel measure ν on X such that

$$p(f) \;=\; \int f \, d\nu$$

for all $f \in C(X)$. Since p properly separates $\mathcal{U}$ and $\mathcal{N}$, it must be that $\nu(X) > 0$. Furthermore, $\nu(X) < \infty$, since otherwise the function that is identically one would not be integrable. Letting $\hat{\sigma} = \frac{1}{\nu(X)}\nu$, I claim that $(\hat{\sigma}, \hat{\sigma})$ is a mixed strategy equilibrium. By construction, for all $\sigma' \in \Sigma^c$, we have

$$U(\hat{\sigma}, \sigma') \;=\; \int f'(x)\hat{\sigma}(dx) \;=\; \frac{1}{\nu(X)} p(f') \;\geq\; 0,$$

where $f' = U(\cdot, \sigma')$. By (C1) and symmetry, $U(x, \hat{\sigma}) \leq 0$ for all $x \in X$. Applying symmetry again, we then have

$$U(x, \hat{\sigma}) \;\leq\; 0 \;\leq\; U(\hat{\sigma}, x')$$

for all $x, x' \in X$, establishing the claim and completing the proof.

Combining Theorem 1 with Propositions 1–3 yields the obvious corollaries in terms of the conditions in those propositions. In particular, the applications considered below all proceed by verifying the conditions in Propositions 2 and 3, which imply (C2) and (C3), which imply (C1). Perhaps less obviously, Theorem 1 extends to non-symmetric zero-sum games, if we replace (C1) with the stronger combination of (C2) and (C4).

Corollary 1 *In any two-player, zero-sum game satisfying (C0), (C2), and (C4), there exists a mixed strategy equilibrium.*

For the proof of the corollary, we note that any non-symmetric game can be converted into a symmetric one by first letting each player have strategy set $\hat{X} = X \times X'$, with one player's strategies denoted (x, y) and the other's denoted (x', y'). Here, we endow $\hat{X}$ with the product topology and corresponding Borel field, and we extend the product measure $\hat{\mu} = \mu \times \mu'$ to this field. We then define the "symmetrized" payoff function

$$\hat{u}((x, y), (x', y')) \;=\; u(x, y') - u(x', y),$$

which is Borel measurable. Thus, we essentially let the players play the original game twice, once where the players assume their original roles and once where the players switch roles. Clearly, $((x, y), (x', y'))$ is a pure strategy equilibrium of the symmetrized game if and only if (x, y') and (x', y) are equilibria of the original game.

We extend $\hat{u}$ to mixed strategy pairs $(\hat{\sigma}, \hat{\sigma}')$ as above. It is straightforward to verify that

$$\begin{aligned} \hat{U}(\hat{\sigma}, \hat{\sigma}') &= \int\int \hat{u}((x, y), (x', y'))\hat{\sigma}(d(x, y))\hat{\sigma}'(d(x', y')) \\ &= U(\hat{\sigma}_1, \hat{\sigma}'_2) - U(\hat{\sigma}'_1, \hat{\sigma}_2), \end{aligned}$$

where a subscript 1 denotes the marginal probability measure on X and a subscript 2 denotes the marginal probability measure on X'. Thus, $(\hat{\sigma}, \hat{\sigma}')$ is a mixed strategy equilibrium of the symmetrized game if and only if the corresponding marginals form mixed strategy equilibria of the original game. Note that the symmetrized game satisfies (C0), (C2), and (C4) if the original game does. To see that this is true of (C4), take any $\hat{\sigma}' \in \hat{\Sigma}'$, and any $(x, y) \in \hat{X}$. Then

$$\hat{U}((x, y), \hat{\sigma}') \;=\; U(x, \hat{\sigma}'_2) - U(\hat{\sigma}'_1, y),$$

and, for any $\epsilon > 0$, (C4) in the original game yields $\sigma_1 \in \Sigma^\circ$ and $\sigma_2 \in \Sigma^{\circ\prime}$ such that

$$U(\sigma_1, \hat{\sigma}'_2) \;\geq\; U(x, \hat{\sigma}'_2) - \frac{\epsilon}{2} \quad \text{and} \quad U(\hat{\sigma}'_1, \sigma_2) \;\leq\; U(\hat{\sigma}'_1, y) + \frac{\epsilon}{2}.$$

Defining $\hat{\sigma} = \sigma_1 \times \sigma_2$ as the product measure, we therefore have $\hat{\sigma} \in \hat{\Sigma}^\circ$ and

$$\hat{U}(\hat{\sigma}, \hat{\sigma}') \;\; = \;\; \hat{U}((x, y), \hat{\sigma}') - \epsilon,$$

which delivers (C4) for the symmetrized game. Theorem 1 then yields a mixed strategy equilibrium of the symmetrized game, and the corollary then follows from the preceding observations.

3 Discussion

3.1 Reny's Better Reply Security

Reny (1999) provides general sufficient conditions for existence of mixed strategy equilibrium, and in symmetric, two-player, zero-sum games, our (C1) implies that the mixed extension of the game satisfies his diagonal better reply security condition. For this special class of games, Reny's condition can be defined as follows: the mixed extension satisfies *diagonal better reply security* if for all $\sigma' \in \Sigma$ such that there exists $x \in X$ with $U(x, \sigma') > 0$,

- there exists $\hat{\sigma} \in \Sigma$ and an open set $G \subseteq \Sigma'$ such that $\sigma' \in G$ and $\inf_{\hat{\sigma}' \in G} U(\hat{\sigma}, \hat{\sigma}') > 0$.

That is, if a player can obtain a positive payoff against σ', then that player can obtain a positive payoff against σ', even when the other player's strategy is allowed to vary within some open set.

In fact, (C1) implies Reny's stronger condition of *better reply security,* which is defined as follows. Let

$$\Gamma \;\; = \;\; \{(\sigma, \sigma', u) \in \Sigma \times \Sigma' \times \Re \mid U(\sigma, \sigma') = u\},$$

and let $\overline{\Gamma}$ denote the closure of Γ, where $\Sigma \times \Sigma' \times \Re$ is given the product topology. Then the mixed extension of a symmetric, two-player, zero-sum game satisfies *better reply security* if for all $(\sigma, \sigma', u) \in \overline{\Gamma}$ such that (σ, σ') is not a mixed strategy equilibrium, either

- there exists $\hat{\sigma} \in \Sigma$ and an open set $G \subseteq \Sigma'$ such that $\sigma' \in G$ and $\inf_{\hat{\sigma}' \in G} U(\hat{\sigma}, \hat{\sigma}') > u$, or

- there exists $\hat{\sigma}' \in \Sigma'$ and an open set $G \subseteq \Sigma$ such that $\sigma \in G$ and $\sup_{\hat{\sigma} \in G} U(\hat{\sigma}, \hat{\sigma}') < -u$.

That is, if (σ, σ') is not an equilibrium, then some player can secure a payoff strictly above the closure payoff (either u or $-u$, depending on the player), even when the other player's strategy is allowed to vary within some open set.

To see the claim, assume (C1) holds for a symmetric game. Let (σ, σ', u) belong to $\overline{\Gamma}$, and suppose (σ, σ') is not an equilibrium. Then, by symmetry and interchangability, either (σ, σ) is not an equilibrium or (σ', σ') is not. Without loss of generality, assume the former. Then there exists $x \in X$ such that $U(x, \sigma) > U(\sigma, \sigma) = 0$. By (C1), there exists $\tilde{\sigma} \in \Sigma^c$ such that $U(\tilde{\sigma}, \sigma) > 0$. Thus, $U(\tilde{\sigma}, \cdot) \colon X \to \Re$ is continuous, and there is an open set $G \subseteq \Sigma$ such that $\sigma \in G$ and, for all $\hat{\sigma}' \in G$, $U(\tilde{\sigma}, \hat{\sigma}') > U(\tilde{\sigma}, \sigma)/2 > 0$. By symmetry, for all $\hat{\sigma} \in G$, we have $U(\hat{\sigma}, \tilde{\sigma}) < U(\sigma, \tilde{\sigma})/2 < 0$. If $u \geq 0$, then better reply security is fulfilled at (σ, σ', u) using the second player by setting $\hat{\sigma}' = \tilde{\sigma}$. If $u < 0$ and (σ', σ') is not an equilibrium, then we can argue as above that the first player has a mixed strategy $\hat{\sigma} \in \Sigma^c$ that keeps U non-negative, and so above u, over some open set of the second player's strategies containing σ', as required. Finally, if $u < 0$ and (σ', σ') is an equilibrium, then σ' keeps U non-negative for all strategies of the second player, and we can set $\hat{\sigma} = \sigma'$ to fulfill better reply security.

The next example demonstrates a symmetric, two-player, zero-sum game that satisfies better reply security yet violates (C1), showing that the latter condition is strictly stronger than the former for this class of games.

Example 1 *A Game that Satisfies Better Reply Security but Not (C1).* Let $X = X' = [0, 1]$ denote the strategy space of two players, and we let u denote the payoffs to player 1 and $-u$ the payoffs to player 2, where u is defined by

$$u(x, y) \quad = \quad \begin{cases} 0 & \text{if } x = 0 \text{ or } y = 0 \\ y - x & \text{else.} \end{cases}$$

To see that this game violates (C1), note that for each $x \in (0, 1]$, $U(x, \cdot)$ is discontinuous at $y = 0$. By construction, then, $\sigma \in \Sigma^c$ implies $\sigma(\{0\}) = 1$. This means, however, that $U(\cdot, \sigma')$ is identically zero if $\sigma' \in \Sigma^c$, so (C1) is violated. In contrast, take any $(\sigma, \sigma', u) \in \overline{\Gamma}$ such that (σ, σ') is not a mixed strategy equilibrium, and suppose without loss of generality that $u \leq 0$. If

$u < 0$, then letting $\hat{\sigma}$ be the point mass on zero, better reply security is trivially fulfilled using player 1. Suppose $u = 0$. If $\sigma'((0,1]) > 0$, then choose $\hat{x} > 0$ small enough that

$$\int_{(0,1]} u(\hat{x}, z)\, \sigma'(dz) + \sigma'(\{0\})(-\hat{x}) \quad > \quad 0,$$

and note that $\hat{x}$ yields a positive expected payoff for player 1 over an open set containing σ'. If $\sigma((0,1]) > 0$, then a symmetric argument fulfills better reply security using player 2. In the remaining case, where $\sigma(\{0\}) = \sigma'(\{0\}) = 1$, (σ, σ') is a pure strategy equilibrium, which is ruled out by assumption. Thus, better reply security is satisfied. ∎

It follows from the above that Theorem 1, which is stated for symmetric games, follows directly from Reny's Theorem 3.1 (or from his Theorem 4.1). The proof in the previous section is included for completeness and because the simple geometric approach used there may be of independent interest.

3.2 Dasgupta and Maskin's Conditions

Dasgupta and Maskin (1986) prove existence of mixed strategy equilibrium, focusing on the case in which strategy sets are non-degenerate, compact intervals in the real line. They decompose their other conditions in three ways: the first is that total payoffs are upper semicontinuous; the second is a restriction on the "number" of discontinuities; and the third, "weak lower semicontinuity," is a restriction on the form of discontinuities. Of course, the first of these conditions is trivially satisfied in zero-sum games. The second condition, for the special case of symmetric, two-player, zero-sum games, is that there is a finite number of one-to-one, continuous functions $f^d \colon \Re \to \Re$ such that if u is discontinuous at (x, x'), then $x = f^d(x')$. An implication is that for each given pure strategy x, there is at most a finite number of strategies x' such that u is discontinuous at (x, x'). By Proposition 2, this is sufficient for (C2) and, in fact, strictly stronger than the latter condition.

A symmetric, two-player, zero-sum game satisfies *weak lower semicontinuity* if for all $x \in \text{int}X$, there exists $\alpha \in [0,1]$ such that for all $x' \in X$, we have

$$\alpha \liminf_{z \downarrow x} u(z, x') + (1 - \alpha) \liminf_{z \uparrow x} u(z, x') \quad \geq \quad u(x, x');$$

furthermore, if $x \in \{\min X, \max X\}$, so that one of these one-sided limits is undefined, then the inequality above must hold with only the well-defined limit on the lefthand side. For example, if $x = \max X$, then $\lim_{z \uparrow x} u(z, x') \geq u(x, x')$. To see that this condition implies the condition of Proposition 3, and therefore (C3), suppose $U(x, \sigma') > 0$, where for simplicity we consider only $x \in \text{int}X$. Taking α as in the definition of weak lower semicontinuity, Fatou's lemma implies

$$\begin{aligned} & \alpha \liminf_{z \downarrow x} \int u(z, x')\, \sigma'(dx') + (1-\alpha) \liminf_{z \uparrow x} \int u(z, x')\, \sigma'(dx') \\ \geq & \int [\alpha \liminf_{z \downarrow x} u(z, x') + (1-\alpha) \liminf_{z \uparrow x} u(z, x')]\, \sigma'(dx') \\ \geq & U(x, \sigma') \\ > & 0. \end{aligned}$$

Thus, either there exists $\underline{z} < x$ such that for all $z \in [\underline{z}, x]$, we have

$$U(z, \sigma') \;=\; \int u(z, x')\, \sigma'(dx') \;>\; 0,$$

or there exists $\overline{z} > x$ such that for all $z \in [x, \overline{z}]$, we have

$$U(z, \sigma') \;=\; \int u(z, x')\, \sigma'(dx') \;>\; 0,$$

In either case, the condition of Proposition 3 is satisfied, as claimed.

The next example shows that weak lower semicontinuity is strictly stronger than the condition of Proposition 3, and therefore strictly stronger than (C3), for the class of games considered here in two ways. First, weak lower semicontinuity demands sufficiently high payoffs for strategies *near* x for one player when the other player uses any x', whereas (C3) demands an absolutely continuous mixed strategy that need have no relation to the original strategy x. Second, the measure of "sufficiently high" is the payoff $u(x, x')$, rather than simply exceeding a payoff of zero, as in (C3).

Example 2 *A Game that Satisfies (C3) but Not Weak Lower Semicontinuity.* Let $X = X' = [0, 1]$, and define u by

$$u(x, y) \;=\; \begin{cases} \frac{y-x}{2} & \text{if } [x = 0 \text{ and } y = 1] \text{ or } [x = 1 \text{ and } y = 0] \\ y - x & \text{else.} \end{cases}$$

Note that given strategies x and z such that $0 < z < x$, z weakly dominates x for the first player. Also, given $x = 0$, the payoffs from z close to zero weakly dominate those from x in the limit: for all $y \in X'$, $\liminf_{z \downarrow x} u(z, y) \geq u(0, y)$. If $U(x, \sigma') > 0$, then it follows that every pure strategy in the interval $(0, \epsilon)$ yields a positive expected payoff for sufficiently small enough $\epsilon > 0$. Thus, this game satisfies the condition of Proposition 3. It violates weak lower semicontinuity, however, because $\liminf_{z \uparrow 1} u(z, 0) = -1 < -1/2 = u(1, 0)$. ∎

Dasgupta and Maskin (1986) extend their conditions and results to m-dimensional strategy spaces, assuming compact, convex strategy sets. Of course, convexity is not assumed in Theorem 1. Their restriction on the "number" of discontinuity points is then that there exist a finite number of one-to-one, continuous functions $f^d \colon \Re \to \Re$ such that if u is discontinuous at (x, x'), then $f^d(x_k) = f^d(x'_k)$ for some $k = 1, \ldots, m$, where x_k is the kth coordinate of the vector x. This still implies that for every strategy x, there is at most a finite number of strategies x' such that u is discontinuous at (x, x'), which is stronger than (C2). In the spatial model with three voters, analyzed in Section 4, Dasgupta and Maskin's restriction is violated because the set of such x' may be uncountable, but (C2) is satisfied nevertheless. Weak lower semicontinuity, generalized to multidimensional strategy spaces, also remains stronger than the corresponding condition (C3).

3.3 Examples of Non-Existence

Since (C0) and (C1) are sufficient for existence, those conditions obviously cannot hold in symmetric, two-player, zero-sum games without mixed strategy equilibria. The next classic example is due to Sion and Wolfe (1957).

Example 3 *Non-Existence of Mixed Strategy Equilibrium.* Let $X = X' = [0, 1]$ denote the strategy space of two players, and define u by

$$u(x, y) = \begin{cases} 1 & \text{if } y < x \\ 0 & \text{if } x = y \\ -1 & \text{if } x + \frac{1}{2} > y > x \\ 0 & \text{if } y = x + \frac{1}{2} \\ 1 & \text{if } y > x + \frac{1}{2}. \end{cases}$$

Since this game satisfies (C0) and (C2), yet has no mixed strategy equilibrium, an implication of Corollary 1 is that it violates (C3). To see this, let

the second player use the mixed strategy σ' with equal probability on 1 and 1/2, i.e., $\sigma'(\{1\}) = \sigma'(\{1/2\}) = 1/2$. Consider the pure strategy $x = 0$ for the first player, and note that $u(x, \sigma') = 1/2$. Note, however, that

$$\sup_{\sigma \in \Sigma^\circ} u(\sigma, \sigma') \quad = \quad 0,$$

contrary to (C3). ∎

A symmetric example was suggested by Michel Le Breton (personal correspondence).

Example 4 *Non-Existence of Mixed Strategy Equilibrium.* Let $X = X' = [0, 1]$, and define u by

$$u(x, y) \quad = \quad \begin{cases} 1 & \text{if } x > y \text{ and either } x < 1 \text{ or } y > 0 \\ 0 & \text{if } x = y \\ -1 & \text{if } x < y \text{ and } x = 0 \text{ and } y = 1. \end{cases}$$

This defines a symmetric, two-player, zero-sum game with no mixed strategy equilibrium. To see this, consider any symmetric equilibrium. If the equilibrium mixed strategy places positive probability on $[0, 1)$, then either player can obtain a positive expected payoff by moving that mass up arbitrarily close to (but below) 1; and if the equilibrium puts probability one on 1, then either player can obtain a positive expected payoff by choosing 0, a contradiction. Since the game has no symmetric equilibrium, it cannot have any equilibrium. Here, (C1) is violated when $x = 0$ and $\sigma'(\{1\}) = 1$. In this case $U(x, \sigma') = 1$, but no hospitable strategy can put positive probability on 0 or 1, and so, for every hospitable strategy $\hat{\sigma}$, $U(\hat{\sigma}, \sigma') = -1$. ∎

Example 4 is especially relevant, because if $u(x, y) = 1$ is interpreted as x being majority-preferred to y, then this is an example of an electoral model with specific majority-preferences. It shows that without some restrictions on voter preferences (and thereby on majority preferences), we are not guaranteed the existence of an equilibrium. In the next section, standard assumptions will be imposed on voter preferences, including continuity, to preclude the latter example.

4 The Classical Spatial Model with Three Voters

Now let $A \subseteq \Re^d$ be a compact, convex subset of Euclidean space with nonempty interior, the elements of which represent possible policy alternatives. Assume a set $N = \{1, 2, 3\}$ of three voters, where the policy preferences of each voter i are represented by a strictly pseudo-concave utility function $u_i \colon A \to \Re$, i.e., u_i is differentiable and, for all distinct $x, y \in A$ with x interior to A, if $u_i(y) \geq u_i(x)$, then $\nabla u_i(x) \cdot (y - x) > 0$. Note the implication that each u_i is strictly quasi-concave on the interior of A. For simplicity, assume that the boundary points of A are Pareto dominated. These assumptions capture the classical special case in which voters have Euclidean preferences and ideal points interior to A.

Define the strict majority preference relation, P, as xPy if and only if $u_i(x) > u_i(y)$ for at least two voters; define weak majority preference, R, as xRy if and only if not yPx, i.e., $u_i(x) \geq u_i(y)$ for at least two voters; and define majority indifference, I, as xIy if and only if xRy and yRx. For any relation Q, which may range over P, R, or I, we let $Q(x) = \{y \in A \mid yQx\}$ denote the upper section of Q, and we let $Q^{-1}(x) = \{y \in A \mid xQy\}$ denote the lower section of Q. Note that by continuity of voter preferences, P is an open relation, while R and I are closed. Furthermore, if xIy, then there is some voter i such that $u_i(x) = u_i(y)$, implying that $I(x)$ has zero Lebesgue measure. The *majority core* consists of the alternatives x such that for all $y \in A$, xRy.

The classical spatial model of elections posits two candidates who, prior to an election, take positions x and x' in the policy space A. Voters, who are not modelled as players, are assumed to vote for candidates with preferred positions, generating a symmetric, two-player, zero-sum game between the candidates with strategy sets $X = X' = A$, endowed with the relative topology, where $\mu = \mu'$ is Lebesgue measure and payoffs are defined as

$$u(x, x') = \begin{cases} 1 & \text{if } xPx' \\ 0 & \text{if } xIx' \\ -1 & \text{if } x'Px. \end{cases}$$

These payoffs, while consistent with the assumption that each candidate seeks only to win the election, are rather stylized: if, for example, one voter prefers

x to x' and the others are indifferent, then $u(x, x') = 0$. It is not assumed that indifferent voters abstain, in which case the first candidate would win with x and should receive a payoff of 1; and it is not assumed that indifferent voters flip coins, in which case the first candidate would win with probability 3/4. The payoffs specified are standard,[6] however, and are relatively nicely behaved — the alternatives introduce additional discontinuities into the model, making them less attractive as starting points.

Mixed strategies, Borel probability measures on A, are denoted by σ and σ', and u is extended to pairs of mixed strategies as in Section 2.

Theorem 2 *In the classical spatial model of elections with three voters, there exists a mixed strategy equilibrium.*

To prove the theorem, it suffices to verify conditions (C2) and (C3), and, by symmetry, it suffices to check them for the first candidate. First, given any $x \in X$, note that u is discontinuous at (x, z) if and only if $z \in I(x)$, which has Lebesgue measure zero. Thus, by Proposition 2, (C2) is satisfied.

Because existence of a pure strategy equilibrium is well-known when the core is non-empty, we simplify the analysis of (C3) by considering only the case of an empty core. Take any $\sigma' \in \Sigma$ and any $x \in X$ such that $U(x, \sigma') > 0$. Since we have assumed that boundary points of A are Pareto dominated, we focus on the case in which x is interior to A. I will construct, for each $\epsilon > 0$, a non-empty open ball $G^\epsilon \subseteq X$ near x such that for all $z \in G^\epsilon$,

$$U(z, \sigma') \;\geq\; U(x, \sigma') - \epsilon.$$

Define

$$U^P(z, \sigma') \;=\; \int_{P(x) \cup P^{-1}(x)} u(z, y)\, \sigma'(dy)$$

and

$$U^I(z, \sigma') \;=\; \int_{I(x)} u(z, y)\, \sigma'(dy),$$

[6] See McKelvey (1986), for example.

and note that

$$U(z,\sigma') \;=\; U^P(z,\sigma') + U^I(z,\sigma').$$

Given such an open ball, Proposition 3 delivers (C3), as required.

It is straightforward to prove that $U^P(\cdot,\sigma')$ is actually continuous. Indeed, letting χ_Y denote the indicator function of a set $Y \subseteq X$, take any $x \in X$ and any sequence $\{x_n\}$ in X converging to x. Note that $\chi_{P(x_n)\cup P^{-1}(x_n)}$ converges to $\chi_{P(x)\cup P^{-1}(x)}$ almost everywhere-σ': take any $y \in P(x) \cup P^{-1}(x)$. If $y \in P(x)$, i.e., $\chi_{P(x)}(y) = 1$, then, since P is open, we must have $y \in P(x_n)$ for high enough n, so $\chi_{P(x_n)}(y) = 1$ for high enough n. A similar argument holds for $y \in P^{-1}(x)$. Thus, Lebesgue's dominated convergence theorem implies that $U^P(x_n,\sigma') \to U^P(x,\sigma')$, as claimed.

Since $U^P(\cdot,\sigma')$ is continuous at x, we may specify an open ball G around x such that for all $z \in G$,

$$U^P(z,\sigma') \;\geq\; U^P(x,\sigma') - \frac{\epsilon}{2}.$$

If $\sigma'(I(x)) = 0$, then we may set $G^\epsilon = G$, and we are done.

Suppose that $\sigma'(I(x)) > 0$. Note that for all $w \in I(x)$, there exists $i \in N$ such that $u_i(x) = u_i(w)$. Let $\phi\colon I(x) \setminus \{x\} \to N$ be any measurable mapping such that for all $w \in I(x) \setminus \{x\}$, $u_{\phi(w)}(x) = u_{\phi(w)}(w)$, and, for each $i \in N$, let $\alpha_i = \sigma'(\phi^{-1}(i))$. That is, $(\alpha_1,\alpha_2,\alpha_3)$ is a distribution on voters induced by the mixed strategy σ' restricted to the points majority-indifferent to x, where α_i is roughly the probability of a point toward which, compared to x, voter i is indifferent. Here, if several voters are indifferent between x and some w, then we allocate the probability to one of them arbitrarily.

Construct a non-zero vector $p^* \in \Re^d$ as follows. If there exists $p \in \Re^d$ such that $p \cdot \nabla u_i(x) > 0$ for each voter i, then, selecting one such vector, let $p^* = p$. If not, then there is a vector $p \in \Re^d$ such that $p \cdot \nabla u_i(x) > 0$ for at least two voters, say 1 and 2. Otherwise, were there no such vector, we would have either $\nabla u_i(x) = 0$ for two voters or $\nabla u_i(x) = 0$ for one voter and $\nabla u_j(x) = -\nabla u_k(x)$ for the remaining two. In the former case, by strict pseudo-concavity, x is the ideal point of two voters, while in the latter case, Plott's (1967) symmetry condition is satisfied at x. Either way, x is a core point, contradicting our maintained assumption in this proof that the core

is empty. Thus, the claim follows. In case $\alpha_1 + \alpha_2 \geq \alpha_3$, set $p^* = p$. In case $\alpha_3 > \alpha_1 + \alpha_2$, there is some $w \in I(x) \setminus \{x\}$ such that $u_3(w) = u_3(x)$. Since xIw, it must be that $u_i(w) \geq u_i(x)$ for at least one other voter, say voter 2. Then pseudo-concavity implies $(w-x)\cdot\nabla u_2(x) > 0$ and $(w-x)\cdot\nabla u_3(x) > 0$, and we set $p^* = w - x$. This construction yields, along with p^*, two voters, j and k (with the third voter denoted by l), such that $\alpha_j + \alpha_k \geq \alpha_l$, that $p^* \cdot \nabla u_j(x) > 0$, and that $p^* \cdot \nabla u_k(x) > 0$.

For each $\delta > 0$, now define $x^\delta = x + \delta p^*$. By strict pseudo-concavity, for δ small enough, we have $u_j(x^\delta) > u_j(x)$ and $u_k(x^\delta) > u_k(x)$, and so $x^\delta P x$. Furthermore, for all $w \in \phi^{-1}(j)$, we may take δ small enough that $x^\delta P w$. The argument is as follows. We have xIw and $u_j(x) = u_j(w)$ by construction. If $u_k(x) \geq u_k(w)$, then $u_j(x^\delta) > u_j(x) = u_j(w)$ and $u_k(x^\delta) > u_k(x) \geq u_k(w)$ for small enough δ implies $x^\delta P w$. Otherwise, if $u_k(x) < u_k(w)$, then xIw implies either that $u_l(x) > u_l(w)$, in which case continuity yields $u_l(x^\delta) > u_l(w)$ for δ small enough, or that $u_l(x) = u_l(w)$. In the latter case, by strict pseudo-concavity, we have $(w - x) \cdot \nabla u_i(x) > 0$ for each voter i, and then, by construction, we have $p^* \cdot \nabla u_i(x) > 0$ for each voter i. This implies that, for small enough δ and for each voter i, we have $u_i(x^\delta) > u_i(x)$, which yields $u_j(x^\delta) > u_j(x) = u_j(w)$ and $u_l(x^\delta) > u_l(x) = u_l(w)$. We then have $x^\delta P w$ for small enough δ, as claimed. A similar claim holds for all $w \in \phi^{-1}(k)$. Thus,

$$\chi_{\phi^{-1}(j)\setminus P^{-1}(x^\delta)} \to 0 \quad \text{and} \quad \chi_{\phi^{-1}(k)\setminus P^{-1}(x^\delta)} \to 0$$

pointwise as δ goes to zero. Therefore, by Lebesgue's dominated convergence theorem, we can choose δ small enough that $x^\delta P x$, that

$$\sigma'(\phi^{-1}(j) \setminus P^{-1}(x^\delta)) < \frac{\epsilon}{8} \quad \text{and} \quad \sigma'(\phi^{-1}(k) \setminus P^{-1}(x^\delta)) < \frac{\epsilon}{16}$$

and that $x^\delta \in G$.

For all $x, z \in X$ such that zPx, we may write

$$\begin{aligned}
&U^I(z, \sigma') \\
&= \sigma'(\{x\}) + \int_{\phi^{-1}(j)} u(z,y)\, \sigma'(dy) \\
&\quad + \int_{\phi^{-1}(k)} u(z,y)\, \sigma'(dy) + \int_{\phi^{-1}(l)} u(z,y)\, \sigma'(dy) \\
&= \sigma'(\{x\}) + \int_{\phi^{-1}(j)\cap P^{-1}(x^\delta)} u(z,y)\, \sigma'(dy) \\
&\quad + \int_{\phi^{-1}(j)\setminus P^{-1}(x^\delta)} u(z,y)\, \sigma'(dy) + \int_{\phi^{-1}(k)\cap P^{-1}(x^\delta)} u(z,y)\, \sigma'(dy) \\
&\quad + \int_{\phi^{-1}(k)\setminus P^{-1}(x^\delta)} u(z,y)\, \sigma'(dy) + \int_{\phi^{-1}(l)} u(z,y)\, \sigma'(dy).
\end{aligned}$$

Let $\{x_n\}$ be any sequence converging to x^δ, and take any $y \in \phi^{-1}(j) \cap P^{-1}(x^\delta)$, so that $u(x^\delta, y) = 1$. Recalling that $x^\delta P x$ and P is open, we have $x_n P x$ and $u(x_n, y) = 1$ for high enough n. Similar remarks hold for $y \in \phi^{-1}(k) \cap P^{-1}(x^\delta)$. Then, by Lebesgue's dominated convergence theorem,

$$\begin{aligned}
&\liminf_{n\to\infty} U^I(x_n, \sigma') \\
&\geq \sigma'(\phi^{-1}(j) \cap P^{-1}(x^\delta)) + \sigma'(\phi^{-1}(k) \cap P^{-1}(x^\delta)) \\
&\quad -\sigma'(\phi^{-1}(j) \setminus P^{-1}(x^\delta)) - \sigma'(\phi^{-1}(k) \setminus P^{-1}(x^\delta)) - \sigma'(\phi^{-1}(l)) \\
&= \sigma'(\phi^{-1}(j)) + \sigma'(\phi^{-1}(k)) - 2\sigma'(\phi^{-1}(j) \setminus P^{-1}(x^\delta)) \\
&\quad -2\sigma'(\phi^{-1}(k) \setminus P^{-1}(x^\delta)) - \sigma'(\phi^{-1}(l)) \\
&= \alpha_j + \alpha_k - \frac{\epsilon}{4} - \alpha_l \\
&\geq -\frac{\epsilon}{4},
\end{aligned}$$

where the last inequality uses $\alpha_j + \alpha_k \geq \alpha_l$. Therefore, there is an open ball G' around x^δ such that

$$U^I(z, \sigma') \geq -\frac{\epsilon}{2}$$

for all $z \in G'$.

Finally, let G^ϵ be any open ball around x^δ contained in $G \cap G'$. Using

$U^I(x, \sigma') = 0$, we then have

$$\begin{aligned} U(z, \sigma') &= U^P(z, \sigma') + U^I(z, \sigma') \\ &\geq U^P(x, \sigma') + U^I(x, \sigma') - \epsilon \\ &= U(x, \sigma') - \epsilon \end{aligned}$$

for all $z \in G^\epsilon$, as desired.

5 More than Three Voters

Unfortunately, Theorem 1 does not readily yield the existence of mixed strategy equilibria in the classical spatial model more generally, because the key condition (C3) becomes extremely difficult to verify. The way it was verified in the proof of Theorem 2 was, for a given point x and mixed strategy σ', to move away from x an arbitrarily small amount to a point x^δ that (i) is strictly majority-preferred to x and (ii) is strictly majority-preferred to close to half of the points in $I(x)$ in the support of σ'. This approach is depicted in Figure 1, where the points majority-preferred to x form three petals with tips y, z, and w, and the points majority-indifferent to x form the boundary of this set.

Suppose σ' puts at least as much probability mass on the indifference curves of voters 1 and 2 in $I(x)$ (indicated by dark curves) as on voter 3's indifference curve in $I(x)$. Any point in the y-petal will be preferred by voters 1 and 2 to x, and such points close enough to x will be strictly majority-preferred to a set that contains, in the limit, the indifference curves of voters 1 and 2 in $I(x)$, as required. Once the point x^δ is found, it is possible to find a sufficiently small open ball around it consisting of points also satisfying (i) and (ii), and then the uniform distribution on this ball fulfills (C3).

The problem that arises with more than three voters is that the sets of voters indifferent between x and y, between x and z, and between x and w may be pairwise disjoint. Term these voters "xy-indifferent," and so on. Moreover, it may be necessary, in order to deviate from x and win against any one of these alternatives, to move to a point preferred to x by *all* of the corresponding indifferent voters. That is, it may be that the only alternatives close to x that are majority-preferred to z are better than z for all xz-indifferent voters. This contrasts with the three-voter example in Figure 1,

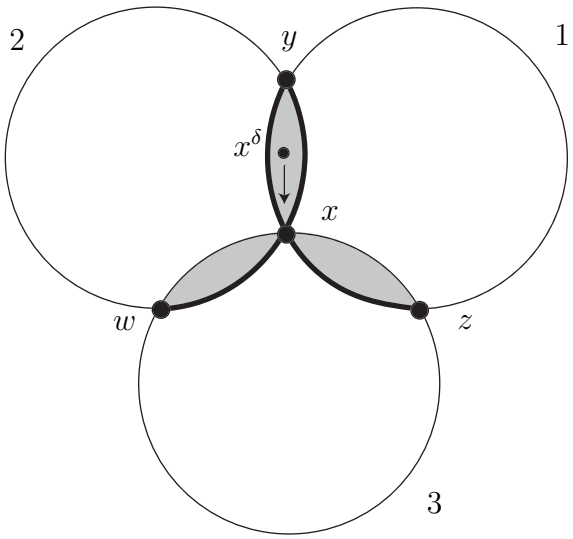

Figure 1: Approach with Three Voters

where voter 1 is indifferent between x and y and between x and z, making it possible to find an alternative x^δ that is strictly majority-preferred to z, despite the fact that one xz-indifferent voter, voter 3, prefers z to x^δ.

Using nine voters, preferences are specified in the following example so that in order to perturb x in a way preferred by all xy-indifferent voters, at least one xz-indifferent voter and at least one xw-indifferent voter must be made worse off. Because all xz-indifferent voters are needed to defeat z, this means that it is impossible to defeat y and z simultaneously. A similar problem is encountered in moving to points preferred to x by all xz-indifferent voters or all xw-indifferent voters.

Example 5 *Condition (C1) is violated, there exists a mixed strategy equilibrium, and diagonal better reply security is satisfied.* Assume A is two-dimensional and nine voters have preferences as depicted in Figure 2. Though preferences are not specified completely here, it is clear that the indifference curves pictured can be generated by differentiable, strictly quasi-concave utility functions. Letting P_i denote the strict preference of voter i, majority preferences are defined as before, namely, xPy if and only if $\#\{i \in N \mid xP_iy\} \geq 5$, and xIy if neither xPy nor yPx. The payoff function of the classical spatial

model is

$$u(x,x') \;=\; \begin{cases} 1 & \text{if } xPx' \\ 0 & \text{if } xIx' \\ -1 & \text{if } x'Px, \end{cases}$$

which again defines a symmetric, two-player, zero-sum game.

The key feature in this construction is the voters' preferences over four points: x, y, z, and w. Preferences restricted to that set are as follows, where parentheses indicate that no restriction over a pair is apparent in Figure 2.

1	2	3	4	5	6	7	8	9
z	z	w	y	(zw)	(yz)	w	y	(yw)
xw	yx	zx	zx	x	x	yx	xw	x
y	w	y	w	y	w	z	z	z

Thus, we have xIy, xIz, and xIw. To flesh out our specification of preferences, assume the parenthetical pairs in the above table are actually indifferences: assume zI_5w, zI_6y, and yI_9w. Thus, y, z, and w are pairwise majority-indifferent as well. Now define

$$\begin{aligned} L_y &= \{v \in A \setminus \{x,y\} \mid vR_ix, i = 2,4,6,7,8,9\} \\ L_z &= \{v \in A \setminus \{x,z\} \mid vR_ix, i = 1,2,3,4,5,6\} \\ L_w &= \{v \in A \setminus \{x,w\} \mid vR_ix, i = 1,3,5,7,8,9\}, \end{aligned}$$

so that L_y, for example, is the y-petal less the base x and tip y. Assume 4 and 6 strictly prefer y to the elements of L_z; assume 1 and 5 strictly prefer z to L_w; and assume 7 and 9 prefer w to L_y. This means that y is majority-preferred to every alternative in the shaded z- and w-petals, excluding x and z and w themselves: for example, voters 4, 6, 7, 8, and 9 strictly prefer y to L_z. Finally, suppose that x, y, z, and w are all strictly majority-preferred to a, as would be the case in Figure 2 if a were moved a sufficient distance from x to the left.

Note that three voters, $\{1,3,5\}$, prefer x to y and four voters, $\{4,6,8,9\}$, prefer y to x, with two, $\{2,7\}$, indifferent. In order to find points arbitrarily close to x that are strictly majority-preferred to y, it is necessary to make both xy-indifferent voters, 2 and 7, better off, meaning that such points must be in the shaded y-petal in Figure 2. Similarly, three voters, $\{7,8,9\}$, prefer

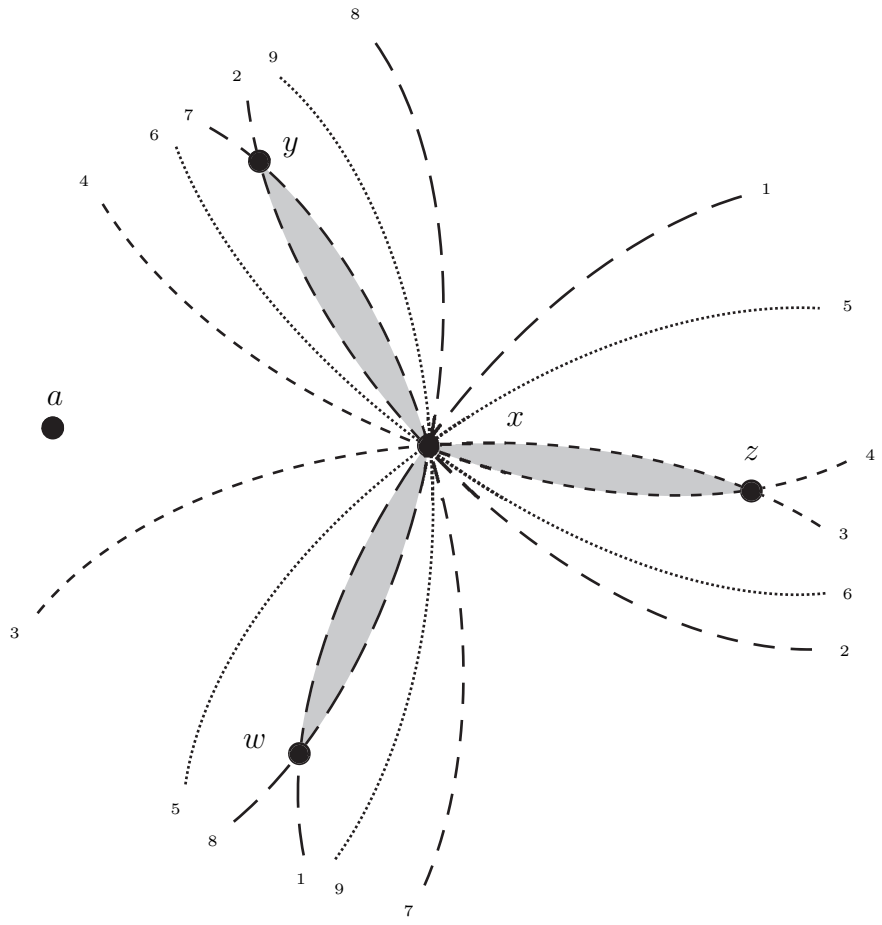

Figure 2: Difficult Case with Nine Voters

x to z and four voters, $\{1, 2, 5, 6\}$, prefer z to x, with two, $\{3, 4\}$, indifferent. To find points arbitrarily close to x that are strictly majority-preferred to z, it is necessary to make both xz-indifferent voters, 3 and 4, better off, meaning that such points must be in the shaded z-petal. And finally, three voters, $\{2, 4, 6\}$, prefer x to w and four voters, $\{3, 5, 7, 9\}$, prefer w to x, with two, $\{1, 8\}$, indifferent. To find points arbitrarily close to x that are strictly majority-preferred to w, it is necessary to make both xw-indifferent voters, 1 and 8, better off, meaning that such points must be in the shaded w-petal. Because the sets of xy-, xz-, and xw-indifferent voters are pairwise disjoint, it will not be possible to find a point close to x that is strictly majority-preferred to two of y, z, and w.

Now suppose that the first candidate takes position x and the second uses the mixed strategy σ' that puts small probability $\epsilon > 0$ on a and puts equal weight on y, z, and w, i.e., $\sigma'(\{y\}) = \sigma'(\{z\}) = \sigma'(\{w\}) = (1 - \epsilon)/3$. Then $U(x, \sigma') = \epsilon > 0$. The approach to verifying (C3) used in the previous section is to find a small open ball around x^δ, arbitrarily close to x, consisting of points that defeat the second candidate's positions most of the time. In order to do that in this example, it is necessary to find points arbitrarily close to x that are strictly majority-preferred to at least two of y, z, and w. Such points must be preferred to x by at least two sets of indifferent voters, which, as argued above, is impossible, and we conclude that (C3) cannot be verified in this way. Can (C3) be verified in other ways? If we let A consist exactly of the shaded petals in Figure 2 together with a,[7] then no. In fact, even the weaker condition (C1) is violated in this example. Against σ', every pure strategy in $L_y \cup L_z \cup L_w$ generates a payoff of $(4\epsilon - 1)/3 < 0$, so the only strategies that achieve a positive payoff against σ' have support in $\{x, y, z, w\}$. None of these strategies are hospitable, nor are mixtures over them.

In fact, there is a mixed strategy equilibrium in this model: let each candidate use the mixed strategy σ^* that puts probability $1/3$ each on y, z, and w. As in Example 1, (C1) is violated, so it is clearly not necessary for existence. Furthermore, we can verify that Reny's (1999) diagonal better reply security is satisfied. Note that the expected payoff of the first candidate from using σ^* against any alternative in $A \setminus \{x, y, z, w\}$ is positive. Now take

[7] This set is non-convex, of course, but that is not usually an issue for existence of mixed strategy equilibria.

any mixed strategy σ', and suppose that there is some alternative v such that $U(v, \sigma') > 0$. If σ' puts positive probability on $A \setminus \{x, y, z, w\}$, then σ^* fulfills the requirement of diagonal better reply security. Otherwise, σ' has support on $\{x, y, z, w\}$, and v must belong to $A \setminus \{x, y, z, w, a\}$. Without loss of generality, suppose v lies in L_y, so that $\sigma'(\{x, y\}) > 1/2$, and then v itself fulfills the requirement of the condition.

Note that the use of mixed strategies in securing deviation payoffs in Reny's condition is crucial, and it may be that the securing mixed strategies must be far from the initial strategy with the positive payoff. For the mixed strategy σ' that puts probability ϵ on a and $(1-\epsilon)/3$ each on y, z, and w, there is no pure strategy v such that $U(v, \sigma') > 0$ and for which there exists an open set $G \subseteq \Sigma'$ with $\sigma' \in G$ and, for all $\hat{\sigma}' \in G$, $U(v, \hat{\sigma}') > 0$. The reason is that $U(v, \sigma') > 0$ implies $v \in \{x, y, z, w\}$. Say $v = y$. But then any mixed strategy $\hat{\sigma}'$ that moves probability $(1-\epsilon)/3$ in σ' from y into L_y yields $U(v, \hat{\sigma}') < 0$, and such a strategy can be found arbitrarily close to σ'. Furthermore, no mixed strategy close to v will fulfill the requirements of diagonal better reply security, for the alternatives near, but not equal to, v generate negative payoffs. ∎

6 The Spatial Model with a Continuum of Voters

Assume that $A \subseteq \Re^d$ is a compact, convex subset of Euclidean space with nonempty interior, and assume a measure space $(\Omega, \mathcal{C}, \nu)$ of voters, where ν is a probability measure on the coalitions C in $\mathcal{C}$. Assume that each voter ω is endowed with a strict preference relation P_ω, with corresponding weak preference relation R_ω and indifference relation I_ω. As usual, we assume P_ω is asymmetric, but for reasons that will become clear in the next section, we depart from Kramer's (1978) model by dropping transitivity properties on voter preferences. Endowing A with the relative topology, assume that P_ω is *continuous,* i.e., P_ω is open as a subset of $A \times A$, and *star-shaped,* i.e., for all $x, y \in A$ and all $\alpha \in (0, 1]$, $xP_\omega y$ implies $\alpha x + (1-\alpha)yP_\omega x$. This latter condition is weaker than the usual convexity assumption and will be exploited in the next section.

Let $P(x,y) = \{\omega \in \Omega \mid xP_\omega y\}$ be the set of voters who strictly prefer x to y, and assume that for all $x, y \in A$, $P(x,y) \in \mathcal{C}$. Let

$$\overline{P}(x,y) \;=\; \{\omega \in \Omega \mid (x,y) \in \text{clos}P_\omega\}$$

be the set of voters who strictly prefer alternatives near x to alternatives near y. When voter indifference curves are "thin," $\overline{P}(x,y)$ is simply the set of voters who weakly prefer x to y. Assume that for all $x, y \in A$, $\overline{P}(x,y) \in \mathcal{C}$. Last, assume that the set of voters who are barely indifferent between any two distinct alternatives is small: for all distinct $x, y \in A$, $\nu(\overline{P}(x,y) \backslash P(x,y)) = 0$. The latter assumption implies that the set Ω is uncountably infinite, so the focus here is on large electorates.

Consider two candidates who simultaneously take positions in the policy space A, and in contrast to the classical spatial model, assume that candidates seek to maximize vote share, rather than probability of winning. Thus, we consider the symmetric, two-player, zero-sum game with strategy sets $X = X' = A$ endowed with the relative topology, where $\mu = \mu'$ is Lebesgue measure, and payoffs defined as

$$u(x,x') \;=\; \nu(P(x,x')) - \nu(P(x',x)).$$

This alternative formulation of payoffs also captures the idea of office-motivated candidates and, in the model with a continuum of voters, it eliminates many discontinuities of the spatial model.

Lemma 1 *In the spatial model with a continuum of voters, for all distinct $x, y \in A$, u is continuous at (x,y).*

To prove the lemma, take distinct $x, y \in A$, and let $(x^n, y^n) \to (x,y)$. By continuity of P_ω, we have

$$P(x,y) \;\subseteq\; \liminf P(x^n,y^n) \quad \text{and} \quad \limsup P(x^n,y^n) \;\subseteq\; \overline{P}(x,y).$$

By assumption, $\nu(P(x,y)) = \nu(\overline{P}(x,y))$. Define

$$C^n \;=\; \bigcap_{m=n}^{\infty} P(x^m,y^m) \quad \text{and} \quad D^n \;=\; \bigcup_{m=n}^{\infty} P(x^m,y^m),$$

and note that $C^n \uparrow \liminf P(x^n, y^n)$ and $D^n \downarrow \limsup P(x^n, y^n)$. Therefore,

$$\nu(P(x,y)) \;\leq\; \nu(\liminf P(x^n, y^n)) \;=\; \nu(\lim C^n) \;=\; \lim \nu(C^n)$$

and

$$\lim \nu(D^n) \;=\; \nu(\lim D^n) \;=\; \nu(\limsup P(x^n, y^n)) \;\leq\; \nu(\overline{P}(x,y)),$$

which implies $\lim \nu(C^n) = \nu(P(x,y)) = \lim \nu(D^n)$. Finally, since $C^n \subseteq P(x^n, y^n) \subseteq D^n$ for all n, we have $\nu(P(x^n, y^n)) \to \nu(P(x,y))$, which implies that u is continuous at (x,y) and proves the lemma.

Theorem 3 *In Kramer's model of elections, there exists a mixed strategy equilibrium.*

The result is clearly true if there exists an strategy $x^* \in X$ such that $\inf_{x' \in X'} u(x^*, x') = 0$. This condition, which is equivalent to $\nu(P(x, x^*)) \leq 1/2$ for all $x \in X$, implies that there is no alternative majority-preferred to x^*, and it makes (x^*, x^*) a pure strategy equilibrium. Thus, assume that for all $x \in X$, $\inf_{x' \in X} u(x, x') < 0$, or equivalently, $\sup_{x' \in X} u(x', x) > 0$.

It suffices to verify (C2) and (C3), and, by symmetry, it suffices to check them for the first candidate. To establish (C2), take any $x \in X$, note that by Lemma 1, we have $\{z \in X \mid u(\cdot, z) \text{ is discontinuous at } x\} \subseteq \{x\}$, which has zero Lebesgue measure. Thus, by Proposition 2, (C2) is satisfied.

For (C3), take any σ' and any $x \in X$ such that $U(x, \sigma') > 0$, and let $y \in X$ satisfy $u(y, x) > 0$, or equivalently, $\nu(P(y, x)) > 1/2$. For each $\alpha \in [0,1]$, define $y(\alpha) = (1-\alpha)x + \alpha y$. By star-shapedness of voter preferences, we have $\nu(P(y(\alpha), x)) > 1/2$, or equivalently, $u(y(\alpha), x) > 0$, for all $\alpha \in (0,1]$. Note that

$$\begin{aligned} U(y(\alpha), \sigma') \;&=\; \int_{X \setminus \{x\}} u(y(\alpha), x') \sigma'(dx') + \sigma'(\{x\}) u(y(\alpha), x) \\ &\geq\; \int_{X \setminus \{x\}} u(y(\alpha), x') \sigma'(dx'), \end{aligned}$$

for all $\alpha \in [0,1]$. By Lemma 1 and Lebesgue's dominated convergence theorem, we have

$$\lim_{\alpha \to 0} U(y(\alpha), \sigma') \;=\; \int_{X \setminus \{x\}} u(x, x') \sigma'(dx') \;=\; U(x, \sigma') \;>\; 0.$$

Thus, there exists $\overline{\alpha} > 0$ such that for all $\alpha \in (0, \overline{\alpha})$, $U(y(\alpha), \sigma') > 0$. Choose $\alpha > 0$ so that $U(y(\alpha), \sigma') > 0$ and, further, that $\sigma'(\{y(\alpha)\}) = 0$, which is possible since the mass points of σ' must be countable. By Lemma 1 and Lebesgue's dominated convergence theorem, $U(\cdot, \sigma')$ is continuous at $y(\alpha)$, so there is an open ball $G \subseteq X$ such that $y(\alpha) \in X$ and, for all $y \in G$, $U(y, \sigma') > 0$. Thus, by Proposition 3, (C3) is satisfied.

7 The Spatial Model with Probabilistic Voting

Again, let $A \subseteq \Re^d$ be a compact, convex subset of Euclidean space with nonempty interior, and now let N be a finite set of voters. Let $(\Theta, \mathcal{T}, \xi)$ be a measure space such that elements θ determine the preferences of voters: let P_i^θ denote the strict preference relation of voter i at θ, with corresponding weak preference relation R_i^θ and indifference relation I_i^θ. Assume P_i^θ is asymmetric, and assume that each P_i^θ is continuous and star-shaped, in the sense of the previous section. The strict majority preference relation at θ, denoted P_θ, is defined as follows: $xP_\theta y$ if and only if

$$\#\{i \in N \mid xP_i^\theta y\} \;>\; \frac{\#N}{2}.$$

It is straightforward to show that P_θ inherits the above properties of voter preferences, so that P_θ is continuous and star-shaped.

The probability measure ξ represents the distribution of voter types. Let $P_i(x, y) = \{\theta \in \Theta \mid xP_i^\theta y\}$ be the type profiles such that voter i strictly prefers x to y, and let

$$\overline{P}_i(x, y) \;=\; \{\theta \in \Theta \mid (x, y) \in \mathrm{clos} P_i^\theta\}$$

be the set of type profiles such that voter i strictly prefers alternatives near x to alternatives near y. Assume that these sets are measurable, and that information about voter preferences is "diffuse," in the sense that for all $i \in N$ and all distinct $x, y \in A$, $\xi(\{\theta \in \Theta \mid \overline{P}_i(x, y) \setminus P_i(x, y)\}) = 0$. That is, we assume that the probability any voter is barely indifferent between distinct alternatives is zero. Consistent with these conventions, let $P(x, y) =$

$\{\theta \in \Theta \mid xP_\theta y\}$ be the type profiles such that x is strictly majority-preferred to y, and let

$$\overline{P}(x,y) \;=\; \{\theta \in \Theta \mid (x,y) \in \mathrm{clos} P_\theta\}$$

be the set of type profiles such that alternatives near x are strictly majority-preferred to alternatives near y.

Consider two candidates who simultaneously take positions in the policy space A, and assume that both candidates seek to maximize the probability of winning. Thus, we consider the symmetric, two-player, zero-sum game with strategy sets $X = X' = A$ endowed with the relative topology, where $\mu = \mu'$ is Lebesgue measure, and payoffs defined as

$$u(x,x') \;=\; \xi(P(x,y)) - \xi(P(y,x)).$$

The next lemma connects the spatial model with probabilistic voting to the framework of the previous section.

Lemma 2 *In the spatial model with probabilistic voting, for all distinct $x, y \in A$, $P(x,y)$ and $\overline{P}(x,y)$ are measurable, and $\xi(\overline{P}(x,y) \setminus P(x,y)) = 0$.*

To prove the lemma, let $\mathcal{M}$ consist of all coalitions of voters containing strictly more than $\#N/2$ members, and note that

$$P(x,y) \;=\; \bigcup_{C \in \mathcal{M}} \bigcap_{i \in C} P_i(x,y),$$

and similarly for $\overline{P}(x,y)$, so the sets are measurable. It is then sufficient to show that

$$\overline{P}(x,y) \setminus P(x,y) \;\subseteq\; \bigcup_{i \in N} [\overline{P}_i(x,y) \setminus P_i(x,y)].$$

To this end, take any $\theta \in \overline{P}(x,y) \setminus P(x,y)$, and let $(x^n, y^n) \to (x,y)$ satisfy $x^n P_\theta y^n$ for all n. For each n, let $C^n \subseteq N$ be such that $\#C^n > \#N/2$ and, for all $i \in C^n$, $x^n P_i^\theta y^n$. Because N is finite, we may without loss of generality go to a subsequence (still indexed by n) such that $C^n = C$ for all n. Clearly, $\theta \in \overline{P}_i(x,y)$ for all $i \in C$. Since it is not the case that $xP_\theta y$, there must exist $i \in C$ such that $\theta \notin P_i(x,y)$, completing the proof of the lemma.

Theorem 4 *In the spatial model with probabilistic voting, there exists a mixed strategy equilibrium.*

The result follows directly from Theorem 3 upon noting that the probabilistic voting model, once majority preferences have been generated, is formally a special case of the general continuum model, as defined the previous section: set $(\Omega, \mathcal{C}, \nu) = (\Theta, \mathcal{T}, \xi)$, and observe that the strict majority preference P_θ fulfills all of the assumptions imposed on voter preferences P_ω above. It is for this reason that in formulating Kramer's model, we omitted transitivity and imposed only the weak convexity condition of star-shapedness. Majority preferences possess desirable transitivity properties only in exceptional circumstances, and even if we assumed each voter's preference P_i^θ were strictly convex in the probabilistic voting model, this property would not generally carry over to the majority preference relation. It was, therefore, out of necessity that we used star-shapedness in the previous section.

References

[1] C. Aliprantis and K. Border (1999) *Infinite Dimensional Analysis: A Hitchhiker's Guide, 2nd ed.,* Springer-Verlag: New York.

[2] E. Aragones and T. Palfrey (2002) "Mixed Equilibrium in a Downsian Model with a Favored Candidate," *Journal of Economic Theory,* 103: 131-161.

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

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

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

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

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

[8] E. Borel (1953) "The Theory of Play and Integral Equations with Skew Symmetric Kernels," *Econometrica,* 21: 97-100.

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

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

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

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

[13] 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.

[14] O. Gross and R. Wagner (1950) "A Continuous Colonel Blotto Game," RAND Research Memorandum, RM-408.

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

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

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

[18] J.-F. Laslier and N. Picard (2002) "Distributive Politics and Electoral Competition," *Journal of Economic Theory,* 103: 106-130.

[19] J.-F. Mertens (1986) "The Minmax Theorem for U.S.C.-L.S.C. Payoff Functions," *International Journal of Game Theory,* 15: 237-250.

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

[21] C. Plott (1967) "A Notion of Equilibrium and its Possibility under Majority Rule," *American Economic Review,* 57: 787-806.

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

[23] W. Rudin (1987) *Real and Complex Analysis, 3rd ed.,* WCB/McGraw Hill: New York.

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

[25] N. Schofield (1983) "Generic Instability of Majority Rule," *Review of Economic Studies,* 50: 695-705.

[26] L. Simon (1987) "Games with Discontinuous Payoffs," *Review of Economic Studies,* 54: 569-597.

[27] M. Sion (1958) "On General Minimax Theorems," *Pacific Journal of Mathematics,* 8: 171-176.

[28] M. Sion and P. Wolfe (1957) "On a Game without a Value," *Contributions to the Theory of Games, III,* Princeton: Annals of Mathematical Studies No 39, 299-306.