Directional Equilibria*

Hun Chung
Dept. of Political Science
University of Rochester

John Duggan
Dept. of Political Science
and Dept. of Economics
University of Rochester

April 7, 2017

Abstract

We propose the solution concept of directional equilibrium for the multidimensional model of voting with general spatial preferences. This concept isolates alternatives that are stable with respect to forces applied by all voters in the directions of their gradients, and it extends a widely (but not well-) known concept from statistics for Euclidean preferences. We establish connections to the majority core, Pareto optimality, and existence and closed graph, and we provide non-cooperative foundations in terms of a local contest game played by voters.

1 Introduction

The multidimensional spatial model of politics provides an abstract framework for collective choice, where points in a Euclidean space can represent vectors of positions on different policy issues. Given an odd number of voters with single-peaked preferences over a single policy issue, the median voter theorem dictates that the median voter's ideal point is the unique element of the majority core. An impediment to the analysis of the general model is the instability of majority rule in multiple dimensions, formalized in the symmetry conditions of Plott (1967) and a result of Schofield (1983) demonstrating the generic emptiness of

*We are grateful for very helpful feedback from Chris Chambers and Bill Zwicker.

the majority core. In reaction to the indeterminacy of majority rule, several solution concepts have discerned alternatives (or sets of alternatives) with special properties as having positive or normative significance. We propose a concept of directional equilibrium that isolates alternatives based on their stability with respect to "forces" applied by voters in the directions of their gradients. This solution generalizes a widely (but not well-) known concept from statistics for Euclidean preferences, and it possesses desirable core consistency, existence, efficiency, and stability properties.

To convey the idea of directional equilibrium, we first consider the problem of maximizing the sum of strictly concave voter utilities over alternatives $x \in \mathbb{R}^m$ to obtain the *utilitarian equilibrium,* which is the unique solution to the first order condition

$$\sum_{i=1}^{n} \nabla u_i(x) \quad = \quad 0. \tag{1}$$

One interpretation of the above first order condition is that the utilitarian equilibrium is stable: if each voter applies to it a force equal to her gradient, then those forces cancel each other out. But this definition of stability is sensitive to scalings of voter utilities and implicitly relies on an interpersonal comparison of utilities. We seek a criterion that is free of interpersonal utility comparisons—in terms of the tug of war analogy, we want to assume that all voters pull with equal force—and we therefore consider the *normalized* gradients of the voters. Thus, for an alternative x that is not an ideal point of any voter, we say x is a *directional equilibrium* if it solves the system

$$\sum_{i=1}^{n} \frac{1}{\|\nabla u_i(x)\|} \nabla u_i(x) \quad = \quad 0 \tag{2}$$

of equations of m equations in m unknowns. If the alternative is the ideal point of a voter, then we assume that voter can resist a net pull in any direction, up to a force of one unit.

This formulation has several desirable properties. We first establish that the concept of directional equilibrium extends the majority core, in the sense that if the number of voters is odd and the majority core is nonempty, then the core alternative is a directional equilibrium. The proof of this result is an immediate corollary of Plott's theorem, for a core alternative must satisfy radial symmetry, which means that each voter is balanced by a voter whose gradient pulls in the opposite direction. In general, this core extension result allows for the possibility of directional equilibria in addition to the core, but we prove that if the set of alternatives is one-dimensional, then the equivalence is exact: if the majority core is nonempty, then the core alternative is the unique directional equilibrium. We then verify that directional equilibria are Pareto optimal and exist in great generality, and that when voter preferences are continuously parameterized, the directional equilibrium correspondence has closed graph. These attributes are appealing for any solution to the spatial model, and existence in particular opens potential for applicability of the concept. Finally, we explore the tug-of-war analogy in greater detail, expressing it in terms of a "contest game" at status quo alternative x, where voters' strategies consist of the application of an amount of force to move the outcome from x. We prove that under weak background conditions, an alternative x is a directional equilibrium if and only if it is the equilibrium outcome of a contest game at x, i.e., the status quo of x maintained in equilibrium, providing non-cooperative foundations for the solution concept.

The directional equilibria generalize a notion of multidimensional median from the statistics literature that has been applied to the spatial model when voter preferences are Euclidean and utilities are linear in distance from a voter's ideal point, i.e., $u_i(x) = -\|\hat{x}^i - x\|$. In this setting, because the norm of a voter's gradient is independent of the alternative $x \neq \hat{x}^i$ at which it is evaluated, the distinction between equations (1) and (2) disappears: an alternative maximizes the sum of voter utilities if and only if the sum of normalized gradients is equal to zero.[1] Put differently, in this special case, the concepts of utilitarian equilibrium and directional equilibrium coincide. Moreover, if the number of voters is odd or the voters' ideal points are not collinear, then it is known that there is a unique utilitarian equilibrium (cf. Baranchuk and Dybvig (2009)), and thus a unique directional equilibrium. Thus, when utilities are linear functions of distance to the voters' ideal points, the directional equilibrium is well-understood and consistent with the utilitarian welfare criterion. In the statistical context, the minimizer of total distance to a given number of points has a long history as a notion of centrality. Weber (1909) introduced the idea in the context of locating a warehouse to minimize transportation costs, and it was imported to the statistics literature by Gini and Galvani (1929) and rediscovered by Haldane (1948). The concept has received various names: mediancentre (Gower (1974)), geometric median (Haldane (1948)), L_1-median (Small (1990)), spatial median (Brown (1983)).[2]

More recently, a literature has developed this concept in the context of social choice theory, but under the assumption of Euclidean preferences. Baranchuk

[1] This observation must be phrased in terms of supergradients if the alternative is the ideal point of some voter, due to non-differentiability of voter utilities.

[2] See Small (1990) for a survey of multidimensional medians.

and Dybvig (2009) assume Euclidean preferences and use the term "consensus," and they apply the concept to analyze decision making by a board of directors. Cervone et al. (2012) use the terminology of "mediancentre" and "Fermat-Weber point," and they discuss computational issues and cite earlier work on the topic. Brady and Chambers (2015) use the term "geometric median," and assuming Euclidean preferences and a variable population, they characterize the geometric median rule by a set of desirable axioms. Brady and Chambers (2016) assume three individuals with Euclidean preferences, and they show that the geometric median is the unique rule satisfying Maskin monotonicity, anonymity, and neutrality.

The concept of directional equilibrium has not received explicit attention in the setting of the general spatial model of politics, however. In general, the systems of equations in (1) and (2) become distinct, and the utilitarian and directional equilibria diverge. Existence of directional equilibria becomes a more subtle issue, as they can no longer be obtained as solutions to a well-defined optimization problem; rather, we use a fixed point argument to establish existence. Despite the relative technical complexity of directional equilibrium, an important benefit of this route is that our concept is invariant with respect to smooth transformations of voter utilities with positive derivative, so that it does not rely on interpersonal comparisons. Applications related to directional equilibrium have been explored in a general spatial setting by Benjamin et al. (2013) and Benjamin et al. (2014), who propose a procedure for making marginal policy adjustments based on normalized gradients of consumers and discuss practical matters concerning the estimation of consumer indifference contours. Hylland and Zeckhauser (1980) develop a closely related idea from a mechanism design standpoint. They propose a mechanism that is essentially the local contest game we define in Section 5, and under assumptions stronger than ours, they prove existence and Pareto optimality of an equilibrium of their mechanism. With the first part of our Theorem 7, the result of Hylland and Zeckhauser (1980) can be used to prove existence of a directional equilibrium, but they do not give a full characterization of the equilibrium outcomes of the contest game (our Theorems 7–9) or connections to the majority core (our Theorems 1 and 2).

Several competing solution concepts have been proposed and explored in the spatial framework. Shepsle (1979) defines the notion of structure-induced equilibrium, which isolates alternatives that are stable with respect to majority voting on each dimension separately. This solution concept extends the core in the same way that directional equilibria do; structure-induced equilibria exist generally; and they can be shown to possess the same generic local uniqueness and stability properties. A drawback of this concept is that, unlike directional equilibrium, structure-induced equilibria may (in three or more dimensions) be Pareto inefficient. McKelvey (1986) extends the concept of uncovered set to the spatial model and defines the yolk as a means of bounding the uncovered set. Whereas the yolk is defined only for Euclidean voter preferences, the uncovered set exists generally and is contained among the Pareto optimal alternatives. A drawback of this concept is that it can contain an open set of alternatives, leading to a limited form of indeterminacy and creating difficulties for comparative statics.[3] Grofman et al. (1987) extend the idea of Copeland winner to the spatial model and define the strong point as the alternative whose win set has minimal Lebesgue measure, and Owen and Shapley (1989) prove uniqueness of the strong point assuming Euclidean voter preferences and a two-dimensional set of alternatives. The existence of at least one strong point follows from elementary arguments, and Pareto optimality is straightforward to show. However, uniqueness (local or otherwise) of the strong point is not known in higher dimensions or when voter preferences are non-Euclidean. Wuffle et al. (1989) define the finagle point, and more recently, Van Wesep (2010) defines the defensive optimum in a model with a continuum of voters, both papers using geometric arguments that rely on the assumption of Euclidean voter preferences.

The remainder of the paper is structured as follows. In Section 2, we present the general spatial model and formal definition of directional equilibrium. In Section 3, we establish connections to the core and Pareto optimality of directional equilibria. In Section 4, we prove that directional equilibria exist in general, and that the equilibrium correspondence has desirable continuity property of closed graph. In Section 5, we relate directional equilibria to the stable outcomes of a local contest game. The appendix contains technical details omitted from the text.

2 Model and Definitions

Let $N = \{1, \ldots, n\}$ be a set of voters, and let $X \subseteq \mathbb{R}^m$ be a non-empty set of alternatives, identified with a subset of m-dimensional Euclidean space. Assume that the preferences of voter i are represented by the utility function $u_i \colon X \to \mathbb{R}$, which is assumed to be continuously differentiable. For simplicity, we also assume that for all $i \in N$, there is a unique *ideal point* $\hat{x}^i \in X$ such that for all $x \in X \setminus \{\hat{x}^i\}$, we have $u_i(\hat{x}^i) > u_i(x)$; and we assume that the ideal point of voter i is the unique critical point of her utility function, i.e., for all $x \in X$, we have $\nabla u_i(x) = 0$ if and

[3] See Banks, Duggan, and Le Breton (2006) and Duggan (2013) for further results on the uncovered set in general environments.

only if $x = \hat{x}^i$. The utility function of voter i is *Euclidean* if she prefers alternatives that are closer to her ideal point, i.e., for all $x, y \in X$, we have $u_i(x) > u_i(y)$ if and only if $\|x - \hat{x}^i\| < \|y - \hat{x}^i\|$. Since our analysis is invariant with respect to smooth, monotonic transformations of utility, as discussed below, we impose a quadratic functional form on Euclidean utilities without loss of generality; that is, if u_i is Euclidean, then we assume $u_i(x) = -\|x - \hat{x}^i\|^2$.

We seek to isolate alternatives that are stable when subject to the exertion of force by voters, without restricting the directions of these forces but assuming that each voter pulls with up to one unit of force. We hypothesize that a voter whose gradient at x is non-zero pulls x in the direction in which utility increases at the highest rate, while a voter whose ideal point is equal to x resists the net pull of other voters. That is, the pull exerted by voter i with $x \neq \hat{x}^i$ is

$$p^i(x) \;=\; \frac{1}{\|\nabla u_i(x)\|}\nabla u_i(x),$$

and if $x = \hat{x}^i$ for voter i, then we set $p^i(x) = 0$. When x is not the ideal point of any voter, it remains in place as long as the net force exerted on it is equal to zero, i.e., $\sum_{i=1}^n p^i(x) = 0$. And if $x = \hat{x}^i$ for a single voter, then i can resist a move in any direction, and so x remains in place as long as i can overcome the net force on x. The magnitude of this net force is just $\|\sum_{j:j\neq i} p^j(x)\|$, so our hypothesis is that x is in equilibrium when $\|\sum_{j:j\neq i} p^j(x)\| \leq 1$. Extending this informal story, if x is the ideal point of k voters, then each of those voters can overcome a net force of up to one unit, and we require that $\|\sum_j p^j(x)\| \leq k$.

Formally, an alternative $x \in X$ is a *directional equilibrium* if

$$\left\|\sum_{i=1}^n p^i(x)\right\| \;\leq\; \#\{i \in N : x = \hat{x}^i\},$$

and we let C^*_{DE} denote the set of directional equilibria. Note that, consistent with the above discussion, if x is not the ideal point of any voter, then it is a directional equilibrium if and only if $\|\sum_{i=1}^n p^i(x)\| = 0$, i.e., the sum of normalized gradients is equal to zero. See Figure 1 for two examples of such directional equilibria with five voters. In the example on the left, the normalized gradients are placed symmetrically to form angles of 120 degrees and clearly sum to zero; while the example on the right is asymmetric, but the normalized gradients at x nevertheless sum to zero. Note that because we normalize the gradient of each voter, we implicitly assume that all interested parties exert equal force. An advantage of this is that $p^i(x)$ is invariant with respect to smooth, monotonic transformations of any voter's utility function, and thus the concept of directional equilibrium is essentially ordinal.

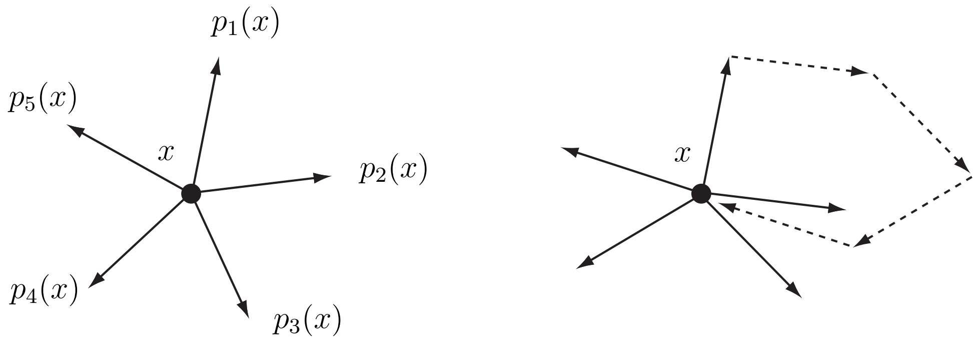

Figure 1: Directional equilibria

To elucidate the concept of directional equilibrium, and to contrast it with Shepsle's (1979) structure-induced equilibrium, we note that the system of equations in (2) depends on the units in which the axes are measured. Given one unit of measure and a solution to (2), we can define an equivalent model by "stretching" the first coordinate axis by the diffeomorphism $f_1 \colon \mathbb{R} \to \mathbb{R}$, where the quantity x_1 is translated as $\tilde{x}_1 = f(x_1)$ in the new model. Then voter i's utility is given by $\tilde{u}_i(\tilde{x}_1, x_{-1}) = u_i(f^{-1}(\tilde{x}_1), x_{-1})$. Directional equilibria in the new model are given by the new system of equations

$$\sum_{i=1}^{n} \frac{1}{\|\nabla \tilde{u}_i(\tilde{x}_1, x_{-i})\|} \nabla \tilde{u}_i(\tilde{x}_1, x_{-i}) \quad = \quad 0,$$

which will generally have solutions that differ from the original.Application of the directional equilibrium concept therefore presumes commitment to units of measurement along the coordinate axes,[4] though a common linear transformation of the units of measure leaves the directional equilibria unchanged. Of course, if all axes are used to measure similar quantities, such as money or distance, then it is natural to use the same unit of measure for all axes, and then the directional equilibria will be invariant to linear transformations of the unit of measure. Note, however, that directional equilibria *are* invariant with respect to rigid Euclidean transformations, such as translation or rotation, of the coordinate axes. In contrast, the structure-induced equilibria are invariant with respect to stretching of

[4]This commitment to units of measure is implicitly made in the large literature on spatial modeling with Euclidean voter preferences. Indeed, many solutions are not even defined for non-Euclidean preferences, so the stretching of an axis will render them inapplicable. This is true more generally in modeling where a specific functional form, e.g., quadratic utility, is assumed.

the axes, as this does not affect the median on any given coordinate. But structure induced equilibria are not generally invariant with respect to rotations of the coordinate axes. These solution concepts each incorporate a different aspect of the Euclidean structure of the spatial model, but an advantage of directional equilibria, established in the next section, is that they are generally Pareto optimal.

3 Core Extension and Pareto Optimality

In this section, we show that the directional equilibria extend the core, in the sense that quite generally, if there is a majority core alternative, then it is necessarily a directional equilibrium. The proof of this follows easily from Plott's (1967) radial symmetry characterization of the majority core. Indeed, let n be odd, consider any majority core alternative $x \in \text{int}X$, and assume there is at most one voter k such that $\nabla u_k(x) = 0$. By Plott's theorem, radial symmetry must be satisfied at x, so there is a permutation $\pi\colon (N \setminus \{k\}) \to (N \setminus \{k\})$ such that for each $i \neq k$, the gradients of i and $\pi(i)$ point in opposite directions. Then

$$\sum_{i:i\neq k} p^i(x) \;=\; \frac{1}{2}\left(\sum_{i:i\neq k}(p^i(x) + p^{\pi(i)}(x))\right) \;=\; 0,$$

which proves the result.

Theorem 1 *Assume that n is odd. Let $x \in intX$, and assume there is at most one voter k such that $x = \hat{x}^k$. If x is a majority core alternative, then it is a directional equilibrium.*

The preceding argument does not preclude the possibility that when the majority core is non-empty, there may be other directional equilibria, but it is straightforward to verify that when the set of alternatives is one-dimensional and voter preferences are single-peaked, the median ideal point is the unique directional equilibrium. Recall that an alternative x is a *median* if $\#\{i \in N : x < \hat{x}^i\} \leq \frac{n}{2}$ and $\#\{i \in N : \hat{x}^i < x\} \leq \frac{n}{2}$. Note that when n is odd, there is a unique median alternative, the ideal point of the median voter.

Theorem 2 *Assume that n is odd, that $X \subseteq \mathbb{R}$ is convex, and that for all $i \in N$, the utility function u_i is strictly quasi-concave. Then the median alternative is the unique directional equilibrium.*

Whereas the preceding strong core equivalence result assumes one dimension and allows quite general preferences, we next specialize to Euclidean voter preferences and allow any number of dimensions. It is known that, quite generally, there is a unique alternative that minimizes total distance to voter ideal points, and that this corresponds with the unique directional equilibrium. The appendix presents a self-contained proof of uniqueness of directional equilibrium when voter preferences are Euclidean, under otherwise quite general conditions. The statement here is from Baranchuk and Dybvig (2009). For general preferences, there may be multiple directional equilibria—even if the number of voters is odd, all voters have strictly quasi-concave utilities, and there is a majority core alternative. Thus, uniqueness does not hold in general.[5]

Theorem 3 (Baranchuk and Dybvig) *Assume X is convex, that voter preferences are Euclidean, and that either n is odd or the voter ideal points $\{\hat{x}^i : i \in N\}$ are not collinear. Then there is at most one directional equilibrium.*

An advantage of directional equilibria is that they are quite generally Pareto optimal; for purposes of comparison, Pareto optimality of structure-induced equilibria is guaranteed only when the set of alternatives is one- or two-dimensional. In addition to differentiability, we assume *strict pseudo-concavity* of voter utilities, i.e., that for all $x \in X$ and all $y \in X \setminus \{x\}$ such that $u_i(y) \geq u_i(x)$, we have $\nabla u_i(x) \cdot (y - x) > 0$. Note that this condition is satisfied if, for example, each u_i is strictly concave.

Theorem 4 *Assume that for all $i \in N$, the utility function u_i is strictly pseudo-concave. If an alternative x is a directional equilibrium, then it is Pareto optimal.*

4 Existence of Directional Equilibrium

This section addresses the issue of existence of directional equilibrium, and it verifies that the equilibrium correspondence has closed graph and, assuming compactness of X, is thereby upper hemicontinuous. The latter property is desirable because it confers a type of robustness of equilibria: if model parameters are varied slightly, it cannot be that directional equilibria are introduced "far" from the equilibrium set for the initial parameterization.

[5]See Figure 3 in the working paper version, Chung and Duggan (2017). We prove there, however, that outside a finitely shy set of voter preferences (see Patty (2007)), the directional equilibria are locally unique, i.e., each equilibrium is unique within a sufficiently small radius.

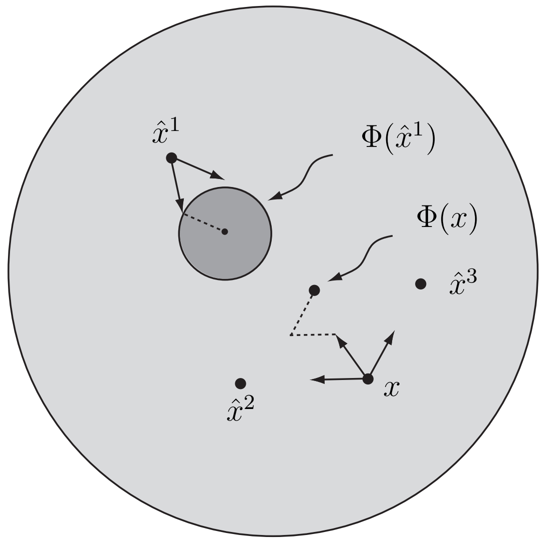

Figure 2: Existence proof

We establish general existence in the next theorem. The proof, which is located in the appendix, is based on an application of Kakutani's fixed point theorem and roughly proceeds by updating any alternative x by adding the sum of normalized gradients, i.e., we map x to $x + \sum_i \epsilon p^i(x)$, suitably scaled by a factor of $\epsilon > 0$. A fixed point of this mapping would of course be a directional equilibrium. But the proof must address two technical difficulties. First, by the assumption that u_i is continuously differentiable, the normalized gradient $p^i(x)$ is continuous at all $x \neq \hat{x}^i$, but it is discontinuous (jumping from norm one to zero) at the ideal point of voter i. To solve this problem, we replace $p^i(\hat{x}^i)$ by a closed disc of radius ϵ, giving us a correspondence, Φ, as depicted in Figure 2. The second difficulty is that the updated alternative may move us outside the set of alternatives, as in Figure 3, where we depict seven voters with $x + \sum_i \epsilon p^i(x) \notin X$. In this case, we scale the move by a factor of $g^\epsilon(x)$ that brings us to the boundary of the set of alternatives.

To facilitate the above proof approach, and with no essential loss of generality, the existence theorem adds two restrictions on the properties of utility functions over the boundary of the set of alternatives. First, we assume that ideal points do not belong to the boundary of the set of alternatives. Second, we assume that the sum of gradients points to the interior of X from any point on the boundary of that set. Under these assumptions, directional equilibria exist generally.

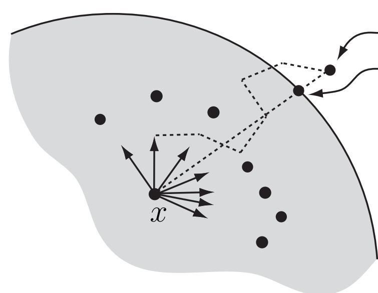

Figure 3: Existence proof again

Theorem 5 *Assume that X is compact and convex, that for all $i \in N$, $\hat{x}^i \notin bdX$, and that for all $x \in bdX$, there exists $\alpha > 0$ such that*

$$x + \alpha \sum_{i=1}^{n} p^i(x) \in intX.$$

Then there is a directional equilibrium, i.e., $C^*_{DE} \neq \emptyset$.

Next, we show that the correspondence of directional equilibria possesses the closed graph property. To formalize this result, let Π be a metric space of parameters, and let utility functions u_i depend on elements $\pi \in \Pi$ of this space parametrically; that is, we now have $u_i \colon X \times \Pi \to \mathbb{R}$. Assume that $u_i(x, \pi)$ is jointly continuous, that for all $\pi \in \Pi$, $u_i(\cdot, \pi)$ is differentiable in x, and that the gradient $\nabla_x u_i(x, \pi)$ is jointly continuous in (x, π). Finally, assume that for each $\pi \in \Pi$, $u_i(\cdot, \pi)$ admits a unique ideal point $\hat{x}^i(\pi)$ and that this is the unique critical point of $u_i(\cdot, \pi)$. We extend our notation in the obvious way: for a voter with non-zero gradient at x, let

$$p^i(x, \pi) \quad = \quad \frac{1}{\|\nabla_x u_i(x, \pi)\|} \nabla_x u_i(x, \pi),$$

and if $\nabla_x u_i(x, \pi) = 0$, then we set $p^i(x, \pi) = 0$, and say an alternative $x \in X$ is a *directional equilibrium at π* if

$$\left\| \sum_{i=1}^{n} p^i(x, \pi) \right\| \quad \leq \quad \#\{i \in N : x = \hat{x}^i(\pi)\},$$

and we let $C^*_{DE}(\pi)$ denote the set of directional equilibria at π. Closed graph of the directional equilibrium correspondence $C^*_{DE}\colon \Pi \rightrightarrows X$ follows straightforwardly from the above definition, and when X is compact, it implies that the correspondence is upper hemicontinuous.

Theorem 6 *The directional equilibrium correspondence $C^*_{DE}\colon \Pi \rightrightarrows X$ has closed graph.*

5 Non-cooperative Stability

The concept of directional equilibrium has direct game-theoretic foundations. Given any alternative $x \in X$ and any $\eta > 0$ such that the closed disc of radius $n\eta$ lies in the set of alternatives, i.e., $D_{n\eta}(x) \subseteq X$, we define the *η-contest game at x* as a strategic form game among the voters such that each voter i's strategy set is the closed disc $D_\eta(0)$ of radius η around zero, and given a profile $s = (s_1, \ldots, s_n)$ of strategies, voter i's payoff is

$$U_i(s) \;=\; u_i\left(x + \sum_j s_j\right).$$

That is, each voter has a "budget" of norm η and can use this budget to pull the outcome in any desired direction. If we interpret x as a status quo outcome, which obtains if the forces applied by the voters cancel each other out, then it is of interest to understand when the status quo is stable, in this sense. Formally, given $\eta > 0$, we say x is *non-cooperatively η-stable* if there is a (pure strategy) Nash equilibrium s in the η-contest game at x such that $x = x + \sum_j s_j$, i.e., $\sum_j s_j = 0$. It is *non-cooperatively strictly-η stable* if there is a strict Nash equilibrium s such that $\sum_j s_j = 0$. Finally, it is *non-cooperatively strongly η-stable* if x is non-cooperatively strictly η-stable, and for every Nash equilibrium s of the η-contest game at x, we have $\sum_j s_j = 0$.

Assuming voter utilities are quasi-concave, the conditions of the Debreu-Fan-Glicksberg theorem are satisfied in the η-contest game at x, and thus there exists a pure strategy equilibrium, say s. The first part of the next theorem establishes that if the voters' strategies cancel each other out, so $\sum_j s_j = 0$, then x must be a directional equilibrium. While quasi-concavity is not actually needed for this direction of implication, it is used to obtain the opposite direction: if utilities are quasi-concave, then every directional equilibrium can be supported as a strict Nash equilibrium outcome of every η-contest game.

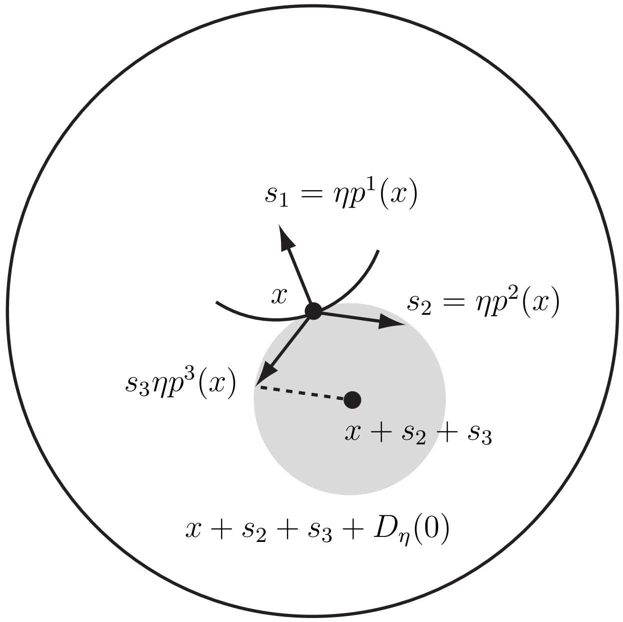

Figure 4: Non-cooperative stability of directional equilibrium

Theorem 7 *For all $x \in X$, if x is non-cooperatively η-stable for some $\eta > 0$ with $D_{n\eta}(x) \subseteq X$, then it is a directional equilibrium. Conversely, assume voter utilities are quasi-concave. If x is a directional equilibrium, then for all $\eta > 0$ with $D_{n\eta}(x) \subseteq X$, x is non-cooperatively strictly η-stable.*

The second part of Theorem 7 is depicted in Figure 4. Here, x is a directional equilibrium, and voter 1's indifference curve through x is shown. Given that the other voters exert effort in the direction of their gradients, i.e., voters 2 and 3 use s_2 and s_3, the set of outcomes that voter 1 can obtain is the shaded disk around $x+s_2+s_3$. By construction, voter 1's indifference curve through x is tangent to the disk; more precisely, the first order condition for voter 1 is satisfied at $s_1 = \eta p^1(x)$. And since the voter's utility function is quasi-concave (with non-zero gradient at x), it follows that x is the unique best outcome achievable for voter 1. The same arguments apply to other voters, and we conclude that $s = (s_1, s_2, s_3)$ is a strict Nash equilibrium of the η-contest game, and that x is indeed non-cooperatively strictly η-stable.

Theorem 7 leaves open the possibility that x is non-cooperatively η-stable, and thereby a directional equilibrium, but that the η-contest game at x admits other Nash equilibria; that is, x may not be non-cooperatively strongly η-stable. This

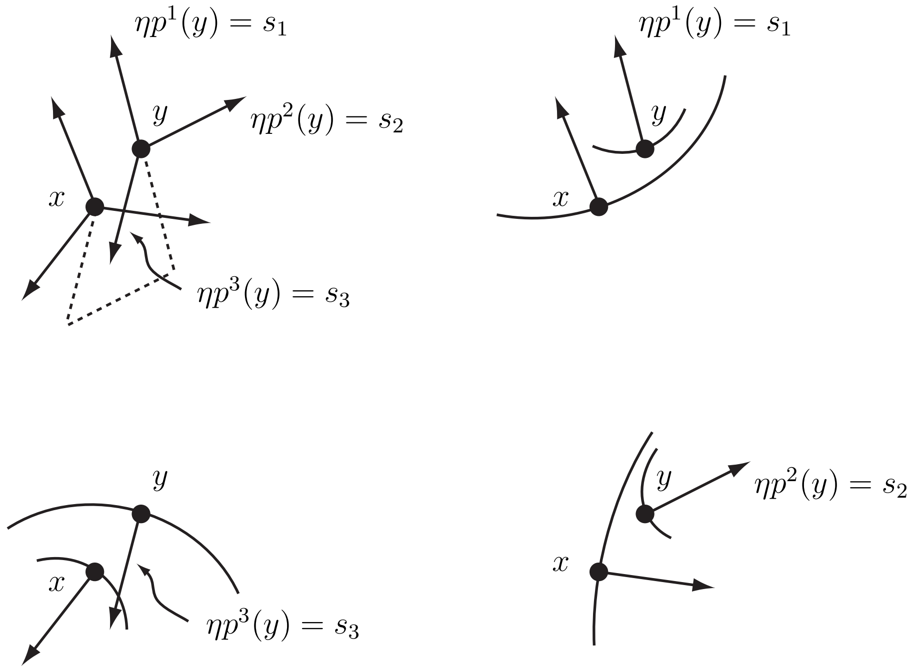

Figure 5: Contest game with multiple equilibria

possibility is illustrated in Figure 5, where x is a directional equilibrium, y is not, but the strategy profile $s = (\eta p^1(y), \eta p^2(y), \eta p^3(y))$ is a Nash equilibrium of the η-contest game at x. To see this in the figure, note that the outcome of this strategy profile is indeed y, and thus (by an argument similar to that for the proof of the first part of Theorem 7) the strategies of the voters are mutual best responses. The equilibrium is depicted in the top-left panel, while the remaining panels depict indifference curves for each voter, showing that the example is consistent with concave utilities. Thus, strong stability is strictly more restrictive than stability.

Nevertheless, the two stability concepts coincide when voter preferences are Euclidean: the properties of Euclidean preferences underpinning the uniqueness result of Theorem 3 can be used here to show that if x is a directional equilibrium, then it is the unique Nash equilibrium of the η-contest game at x. The result in fact uses Lemma 1, but it holds even without the additional background conditions of Theorem 3.

Theorem 8 *Assume voter utilities are Euclidean. For all $x \in X$, if x is a directional equilibrium, then for all $\eta > 0$ with $D_{n\eta}(x) \subseteq X$, x is non-cooperatively strongly η-stable.*

It is clear from Figure 4 that quasi-concavity of voter utilities plays an important role in the second part of Theorem 7. Even without this convex structure, however, we can typically provide local non-cooperative support for directional equilibria. We next consider directional equilibria x that are *non-degenerate,* in the sense that for all voters i, the Hessian $D^2u_i(x)$ of the voter's utility function is negative definite at x. Such non-degenerate alternatives should be expected to be common; indeed, if voter utilities are quadratic, then every alternative is non-degenerate. Substituting non-degeneracy for quasi-concavity, the next result establishes that directional equilibrium implies non-cooperative η-stability for η sufficiently small.

Theorem 9 *For all $x \in intX$, if x is a non-degenerate directional equilibrium, then there exists $\eta > 0$ such that for all $\epsilon \in (0, \eta)$, x is non-cooperatively strictly ϵ-stable.*

As a matter of interest, we remark that quasi-concavity in the second part of Theorem 7 and non-degeneracy in Theorem 9 are needed for the results: it may be that an alternative is a (degenerate) directional equilibrium, yet not non-cooperatively ϵ-stable for any $\epsilon > 0$. For an example, we let $m = 2$ and $n = 3$, and specify utilities as follows:

$$\begin{aligned} u_1(x, y) &= y + x^3 + x^3 \sin(1/x) \\ u_2(x, y) &= -x^2 - y^2 \\ u_3(x, y) &= -x^2 - (1 - y)^2, \end{aligned}$$

where $\sin(1/x)$ and $\cos(1/x)$ are evaluated as zero when $x = 0$. Of course, voters 2 and 3 have Euclidean preferences with ideal points at the origin $(0, 0)$ and at $(0, -1)$, respectively, while voter 1's utility function violates quasi-concavity. Noting that $\frac{d}{dx}x^3 \sin(1/x) \to 0$ as $x \to 0$, voter 1's utility is continuously differentiable, and the gradient of u_1 is

$$\nabla u_1(x, y) = (3x^2 + 3x^2 \sin(1/x) + x \cos(1/x), 1),$$

so evaluated at zero, it is $\nabla u_1(0, 0) = (0, 1)$. Thus, the origin is a directional equilibrium. Note, however, that the Hessian of u_1 is singular, so Theorem 9 cannot be applied. We show in the appendix that this directional equilibrium fails to be non-cooperatively ϵ-stable for any $\epsilon > 0$.

A Technical Details

Proof of Theorem 2: We first verify that the median, say x, is a directional equilibrium. Let $G^{>} = \{i \in N : x < \hat{x}^i\}$ be the voters whose ideal point is above the median, and let $G^{<} = \{i \in N : \hat{x}^i < x\}$ be the voters whose ideal point is below x. Assume without loss of generality that $\#G^{>} \geq \#G^{<}$. Note that for all $i \in G^{>}$ and all $j \in G^{<}$, we have $Du_i(x) > 0$ and $Du_j(x) < 0$, which implies

$$\left\| \sum_{i=1}^{n} p^i(x) \right\| \;=\; \left\| \sum_{i \in G^{>} \cup G^{<}} p^i(x) \right\| \;=\; \#G^{>} - \#G^{<}.$$

Using the identity $\#G^{>} + \#G^{<} + \#\{i \in N : x = \hat{x}^i\} = n$, this yields

$$\left\| \sum_{i=1}^{n} p^i(x) \right\| \;=\; 2\#G^{>} - n + \#\{i \in N : x = \hat{x}^i\}.$$

Since x is the median, we have $\#G^{>} \leq \frac{n}{2}$, i.e., $2\#G^{>} - n \leq 0$, and we conclude from the above that x is a directional equilibrium. Next, assume x is a directional equilibrium. If x is not the median, then using the notational convention above, we can assume without loss of generality that $\#G^{>} > \frac{n}{2}$, but then

$$\begin{aligned} \left\| \sum_{i=1}^{n} p^i(x) \right\| &= \#G^{>} - \#G^{<} \;=\; 2\#G^{>} - n + \#\{i \in N : x = \hat{x}^i\} \\ &> \#\{i \in N : x = \hat{x}^i\}, \end{aligned}$$

a contradiction. □

Our analysis of uniqueness in Theorem 3 relies on the following lemma, which is also used in a key step in the proof of Theorem 8 in our analysis of non-cooperative foundations.

Lemma 1 *Assume X is convex and voter preferences are Euclidean. Given any distinct $x, y \in X$ and any $\alpha \in (0,1)$, let $t = \frac{1}{\|y-x\|}(y - x)$ be the direction pointing from x to y. Then for all voters i, the dot product $p^i((1-\alpha)x + \alpha y) \cdot t$ is non-increasing in α; and if $\{\hat{x}^i, x, y\}$ are not collinear, then it is strictly decreasing in α.*

Proof: Consider any voter i and $\alpha \in (0,1)$. If $\hat{x}^i = (1-\alpha)x + \alpha y$, then the dot product $p^i((1-\alpha)x + \alpha y) \cdot t$ is discontinuous and jumps down at α. Otherwise, if $\hat{x}^i \neq (1-\alpha)x + \alpha y$, then it is differentiable. The first part of the lemma follows

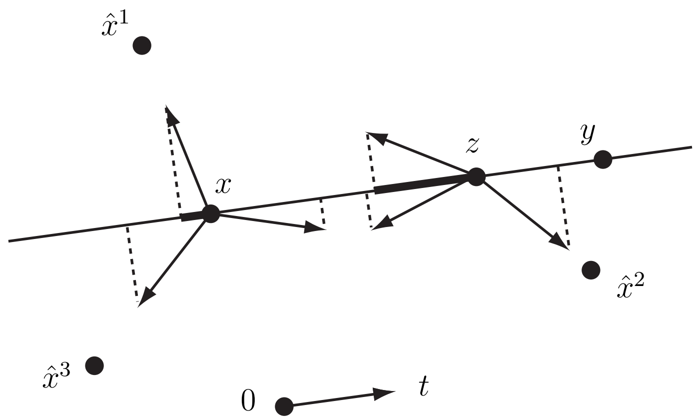

Figure 6: Decreasing dot products

by confirming that this derivative is non-positive. Write $z(\alpha) = (1-\alpha)x + \alpha y$, so that $\nabla z(\alpha) = y - x$. Note that

$$\frac{d}{d\alpha}\frac{(\hat{x}^i - z(\alpha)) \cdot t}{\|\hat{x}^i - z(\alpha)\|} \quad \propto \quad -\|\hat{x}^i - z(\alpha)\|\|y-x\| - [(\hat{x}^i - z(\alpha)) \cdot t]\frac{d}{d\alpha}\|\hat{x}^i - z(\alpha)\|,$$

where $\propto$ indicates that the left-hand side has the same sign as the right-hand side. Here,

$$\frac{d}{d\alpha}\|\hat{x}^i - z(\alpha)\| \quad = \quad \frac{-(\hat{x}^i - z(\alpha)) \cdot (y-x)}{\|\hat{x}^i - z(\alpha)\|} \quad = \quad -\frac{\|y-x\|\,(\hat{x}^i - z(\alpha)) \cdot t}{\|\hat{x}^i - z(\alpha)\|}.$$

Thus, the derivative is non-positive if

$$\|\hat{x}^i - z(\alpha)\|^2 \quad \geq \quad [(\hat{x}^i - z(\alpha)) \cdot t]^2.$$

By the Cauchy-Schwartz inequality, we have $|(\hat{x}^i - z(\alpha)) \cdot t| \leq \|\hat{x}^i - z(\alpha)\|\,\|t\| = \|\hat{x}^i - z(\alpha)\|$, and the desired inequality follows. Moreover, if $\{\hat{x}^i, x, y\}$ are not collinear, then $\{(\hat{x}^i - z(\alpha)), t, 0\}$ are not collinear, and the inequality holds strictly, delivering the second part of the lemma. □

The idea of the lemma is illustrated in Figure 6. Starting with alternative x, the figure indicates the normalized gradients of three voters, which are projected onto the line through x and y. The projection of voter 1's normalized gradient, for example, is the thick line segment emanating from x. Letting t be the direction pointing from x to y, the fact that the projection of $p^1(x)$ onto the line is on the side of x opposite y indicates that the dot product $t \cdot p^i(x)$ is negative. As we move the alternative toward y, voter gradients continue to point toward their ideal point, and these will tip away from y; this is evident in the case of voter 1, because the line segment emanating from z is longer, indicating that the dot product $t \cdot p^i(z)$ is negative and of greater magnitude.

Proof of Theorem 3: In case ideal points are collinear, then n is odd, and since voter preferences are Euclidean, uniqueness follows from Theorem 2, applied to the one-dimensional model induced by voter preferences. In the remaining case that voter ideal points are not collinear, suppose there are distinct directional equilibria, say x and y, and let $t = \frac{1}{\|x-y\|}(y-x)$ be the direction pointing from x to y. Let $G = \{i \in N : \hat{x}^i = x\}$ and $H = \{i \in N : \hat{x}^i = y\}$ be the coalitions of voters with ideal points at x and y, respectively. Note that by the Cauchy-Schwartz inequality, we have

$$t \cdot \sum_{i \in N \setminus G} p^i(x) \;\; \leq \;\; \left\| \sum_{i \in N \setminus G} p^i(x) \right\|.$$

Moreover, for all $\alpha \in (0,1)$ and all $i \in G$, we have $p^i((1-\alpha)x + \alpha y) = -t$, where we use the fact that $(1-\alpha)x + \alpha y = x + \alpha\|y - x\|t$. Thus, we have

$$\lim_{\alpha \downarrow 0} t \cdot \sum_{i=1}^{n} p^i((1-\alpha)x + \alpha y) \;\; \leq \;\; \left\| \sum_{i \in N \setminus G} p^i(x) \right\| - \#G \;\; \leq \;\; 0,$$

where the second inequality follows from the definition of directional equilibrium. A similar argument, starting from y, yields

$$\lim_{\alpha \uparrow 1} t \cdot \sum_{i=1}^{n} p^i((1-\alpha)x + \alpha y) \;\; \geq \;\; \#H - \left\| \sum_{i \in N \setminus H} p^i(x) \right\| \;\; \geq \;\; 0. \tag{3}$$

But starting from x, as we increase α from zero to one, Lemma 1 implies that for each voter i, the quantity $t \cdot p^i((1-\alpha)x + \alpha y)$ weakly decreases; this change is continuous unless $\hat{x}^i = (1-\alpha)x + \alpha y$, in which case the quantity jumps down discontinuously at $\hat{x}^i$. Moreover, by the assumption that ideal points are not collinear, there is at least one voter j whose ideal point does not lie on the line through x and y, and for this voter, Lemma 1 implies that $t \cdot p^j((1-\alpha)x + \alpha y)$ is strictly decreasing in α. We conclude that

$$\lim_{\alpha \uparrow 1} t \cdot \sum_{i=1}^{n} p^i((1-\alpha)x + \alpha y) \;\; < \;\; 0,$$

contradicting (3). □

Proof of Theorem 4: Consider a directional equilibrium x, and let $y \in X \setminus \{x\}$ be any other alternative. We consider two cases. First, if $\nabla u_k(x) = 0$ for some voter k, then it follows that $x = \hat{x}^k$ and $u_k(x) > u_k(y)$, so y does not Pareto dominate x, and we conclude that x is Pareto optimal. Second, suppose $\nabla u_i(x) \neq 0$ for all $i \in N$. For each i, define the quantity

$$\alpha_i \;=\; \frac{1}{\|\nabla u_i(x)\| \sum_{i=1}^{n} \frac{1}{\|\nabla u_i(x)\|}},$$

and note that $\alpha_i > 0$ for each i and $\sum_{i=1}^{n} \alpha_i = 1$. Since x is a directional equilibrium, we have

$$\sum_{i=1}^{n} \alpha_i \nabla u_i(x) \;=\; \frac{1}{\sum_{i=i}^{n} \frac{1}{\|\nabla u_i(x)\|}} \sum_{i=1}^{n} \frac{1}{\|\nabla u_i(x)\|} \cdot \nabla u_i(x) \;=\; \beta \sum_{i=1}^{n} p^i(x) \;=\; 0,$$

where β is a constant. Then we have

$$0 \;=\; (y-x)\cdot 0 \;=\; (y-x)\cdot \sum_{i=1}^{n} \alpha_i \nabla u_i(x) \;=\; \sum_{i=1}^{n} \alpha_i (y-x)\cdot \nabla u_i(x),$$

and this in turn implies that there is a voter i such that $\alpha_i(y-x)\cdot\nabla u_i(x) \leq 0$. And since $\alpha_i > 0$, it follows that $(y-x)\cdot\nabla u_i(x) \leq 0$. Since u_i is strictly pseudo-concave, we conclude that $u_i(x) > u_i(y)$, and once again, y does not Pareto dominate x, so x is Pareto optimal. □

Proof of Theorem 5: Since we assume $\hat{x}^i \in \text{int}X$ for all $i \in N$, there exists an $\epsilon > 0$ such that for all $i \in N$, the closed disc of radius $n\epsilon$ around $\hat{x}^i$ is contained in the interior of the set of alternatives, i.e., $D_{n\epsilon}(\hat{x}^i) \subseteq \text{int}X$. Given any $x \in X$, consider the following constrained maximization problem,

$$\begin{aligned} &\max_{\alpha\in\mathbb{R}} \alpha \\ \text{s.t. } & x + \alpha \textstyle\sum_{i=1}^{n} \epsilon p^i(x) \in X \\ &\alpha \in [0,1]. \end{aligned}$$

Since the constraint set is nonempty (it contains $\alpha = 0$) and compact, this problem has a solution, which must be unique. We denote the solution by $g^\epsilon(x)$. Note that by construction of ϵ, we have $g^\epsilon(x) = 1$ when $x = \hat{x}^i$ for some voter i, and by the assumption that the sum of normalized gradients points to the interior of X, we have $g^\epsilon(x) > 0$.

Claim 1: For all $x \in X$, g^ϵ is continuous at x.

We first show that the constraint correspondence $\Gamma\colon X \rightrightarrows [0,1]$ defined by

$$\Gamma(x) \;=\; \left\{\alpha \in [0,1] : x + \alpha \sum_i \epsilon p^i(x) \in X\right\}$$

is continuous at all $x \in X' = X \setminus \{\hat{x}^1, \ldots, \hat{x}^n\}$. The correspondence obviously has closed graph on X', and since the unit interval is compact, it is upper hemicontinuous. To prove lower hemicontinuity, note that the correspondence $\Gamma^\circ\colon X' \rightrightarrows [0,1]$ defined by

$$\Gamma^\circ(x) \;=\; \left\{\alpha \in [0,1] : x + \alpha \sum_i \epsilon p^i(x) \in \mathrm{int}X\right\}$$

has open graph and is therefore lower hemicontinuous. We claim that for all $x \in \mathrm{int}X'$, we have $\Gamma(x) = \mathrm{clos}\Gamma^\circ(x)$. Indeed, given $\alpha \in \Gamma(x)$, we have $x + \alpha \sum_i \epsilon p^i(x) \in X$. Since X is convex and $x \in \mathrm{int}X$, it follows that for each k,

$$x + \frac{(k-1)\alpha}{k} \sum_i \epsilon p^i(x) \in \mathrm{int}X, \tag{4}$$

so $\frac{(k-1)\alpha}{k} \in \Gamma^\circ(x)$ and $\lim_{k\to\infty} \frac{(k-1)\alpha}{k} = \alpha$, as claimed. Furthermore, we claim that for all $x \in \mathrm{int}X' \setminus \{\hat{x}^1, \ldots, \hat{x}^n\}$, we also have $\Gamma(x) = \mathrm{clos}\Gamma^\circ(x)$. Indeed, now (4) holds by convexity of X and our assumption that the sum of normalized gradients points to the interior of X, and the argument proceeds as above, as claimed. We conclude that for all $x \in X'$, $\Gamma(x)$ is the closure of the value of a lower hemi-continuous correspondence; thus, Γ is lower hemi-continuous on X'. By the theorem of the maximum, it follows that $g^\epsilon\colon X' \to [0,1]$ is continuous. We have left to verify that g^ϵ is continuous at the ideal point $\hat{x}^i$ of a voter. We have noted that $g^\epsilon(\hat{x}^i) = 1$. Moreover, by choice of ϵ, we have $\hat{x}^i + D_{n\epsilon}(\hat{x}^i) \subseteq \mathrm{int}X$. Given any sequence $\{x^k\}$ in X such that $x^k \to \hat{x}^i$, we have

$$\left\| x^k + \sum_j \epsilon p^j(x^k) - \hat{x}^i \right\| \;\le\; \|x^k - \hat{x}^i\| + \sum_j \epsilon \|p^j(x^k)\| \;\le\; \|x^k - \hat{x}^i\| + n\epsilon \;\to\; n\epsilon.$$

Therefore, for high enough k, we have $x^k + \sum_j \epsilon p^j(x^k) \in \mathrm{int}X$, so $g^\epsilon(x^k) = 1$, and we conclude that g^ϵ is continuous. □

Given any $x \in X$, let $G(x) = \{i \in N : \hat{x}^i \neq x\}$ be the set of voters whose ideal point is distinct from x, and let $H(x) = \{i \in N : \hat{x}^i = x\}$ be the set of voters whose ideal point is equal to x. Next, define the correspondence $\Phi : X \rightrightarrows X$ as follows: for all $x \in X$,

$$\Phi(x) = \left[x + g^{\epsilon}(x) \sum_{j \in G(x)} \epsilon p^j(x) + \sum_{j \in H(x)} D_{\epsilon}(0) \right] \cap X,$$

where $D_{\epsilon}(0)$ is the closed disc of radius ϵ around zero, and we let the summation over the empty set equal zero or the singleton consisting of the zero vector, as appropriate.

Claim 2: For all $x \in X$, $\Phi(x)$ is non-empty, convex, and closed.

By construction, $x + g^{\epsilon}(x) \sum_{j \in G(x)} \epsilon p^j(x) \in X$. Since $0 \in D_{\epsilon}(0)$, we have $\Phi(x) \neq \emptyset$. Since $D_{\epsilon}(0)$ is convex and the sum of convex sets is convex, we immediately obtain convexity of $\Phi(x)$. Next, consider any sequence $\{y^k\}$ in $\Phi(x)$ converging to y. In case $H(x) = \emptyset$, we have $y^k = x + g^{\epsilon}(x) \sum_{i=1}^n \epsilon p^i(x)$ for all k, and thus $y = y^k \in \Phi(x)$. Otherwise, in case $H(x) \neq \emptyset$, we use the fact that $D_{\epsilon}(0)$ is compact, and therefore the sum of closed discs is closed. □

Claim 3: The correspondence Φ is upper hemi-continuous.

Consider any sequence $\{x^k\}$ in X converging to $x \in X$, and let $\{y^k\}$ be a sequence such that $y^k \in \Phi(x^k)$ for all k and $y^k \to y$. We must show that $y \in \Phi(x)$. In case $H(x) = \emptyset$, we have $H(x^k) = \emptyset$ for high enough k, which implies

$$y^k = x^k + g^{\epsilon}(x^k) \sum_{i=1}^n \epsilon p^i(x^k) \to x + g^{\epsilon}(x) \sum_{i=1}^n \epsilon p^i(x) = y \in \Phi(x),$$

where the limit uses continuous differentiability of u_i and continuity of g^{ϵ} from Claim 1. In case $H(x) \neq \emptyset$, if $x^k = x$ for sufficiently high k, then the claim follows by closedness of $\Phi(x)$ from Claim 2. Thus, suppose that $x^k \neq x$ for arbitrarily high k, so (going to a subsequence if necessary) we can assume that $H(x^k) = \emptyset$ for all k. Thus, we have

$$y^k = x^k + g^{\epsilon}(x^k) \sum_{j=1}^n \epsilon p^j(x^k) = x^k + g^{\epsilon}(x^k) \sum_{j \in G(x)} \epsilon p^j(x^k) + g^{\epsilon}(x^k) \sum_{j \in H(x)} \epsilon p^j(x^k)$$

for all k. We can assume (going to a subsequence if necessary) that for all $j \in H(x)$, the sequence $\{p^j(x^k)\}$ converges to a limit $w^j \in D_1(0)$. Thus, taking limits, we obtain

$$y = \lim_{k \to \infty} y^k = x + g^{\epsilon}(x) \sum_{j \in G(x)} \epsilon p^j(x) + g^{\epsilon}(x) \sum_{j \in H(x)} \epsilon w^j.$$

Noting that $g^\epsilon(x) = 1$ and $\epsilon w^j \in D_\epsilon(0)$, we then have $y \in \Phi(x)$, as required. □

Thus, Kakutani's fixed point theorem guarantees the existence of a fixed point, i.e., an alternative $x \in X$ such that $x \in \Phi(x)$. The final step of our proof is to show that the fixed point is a directional equilibrium. First, suppose $x \in \Phi(x)$ is not the ideal point of any voter, so $H(x) = \emptyset$. Then

$$x = x + g^\epsilon(x) \sum_{i=1}^{n} \epsilon p^i(x),$$

which implies

$$g^\epsilon(x) \sum_{i=1}^{n} \epsilon p^i(x) = 0.$$

Dividing by $g^\epsilon(x)\epsilon > 0$, we then have

$$\sum_{i=1} p^i(x) = 0,$$

and we conclude that x is a directional equilibrium. Next, suppose $x \in \Phi(x)$ is an ideal point, so that $H(x) \neq \emptyset$. Then for each $j \in H(x)$, there exists a $v^j \in D_\epsilon(0)$ such that

$$x = x + \sum_{j\in G(x)} \epsilon p^j(x) + \sum_{j \in H(x)} v^j,$$

which implies

$$\left\| \sum_{j\in G(x)} \epsilon p^j(x) \right\| = \left\| \sum_{j\in H(x)} v^j \right\| \leq \sum_{j\in H(x)} \|v^j\| \leq \epsilon \, \#H(x).$$

Dividing by $\epsilon > 0$, this implies

$$\left\| \sum_{j=1}^{n} p^j(x) \right\| = \left\| \sum_{j\in G(x)} p^j(x) \right\| \leq \#\{i \in N : x = \hat{x}^i\},$$

and we again conclude that x is a directional equilibrium. □

Proof of Theorem 6: Let $\{x^m\}$ and $\{\pi^m\}$ be sequences in X and in Π, respectively, such that for all m, we have $x^m \in C^*_{DE}(\pi^m)$. If $\nabla_x u_i(x^m, \pi^m) = 0$ for all m, then by continuity we have $\nabla_x u_i(x,\pi) = 0$. Contrapositively, if $\|p^i(x,\pi)\| = 1$, then $\|p^i(x^m,\pi^m)\| = 1$ for large enough m. Combining these observations, we have

$$\begin{aligned} \left\| \sum_{i=1}^n p^i(x,\pi) \right\| &\leq \liminf_{m\to\infty} \left\| \sum_{i=1}^n p^i(x^m,\pi^m) \right\| \\ &\leq \limsup_{m\to\infty} \#\{i \in N : x = \hat{x}^i(\pi^m)\} \\ &\leq \#\{i \in N : x = \hat{x}^i(\pi)\}. \end{aligned}$$

We conclude that $x \in C^*_{DE}(\pi)$, as required. □

Proof of Theorem 7: First, let s be a Nash equilibrium of the η-contest game at x with $\sum_j s_j = 0$, so for all voters i, the strategy s_i solves

$$\begin{array}{c} \max_{y\in\mathbb{R}^m} u_i\left(x + (\sum_{j:j\neq i} s_j) + y\right) \\ \text{s.t. } \|y\|^2 \leq \eta^2. \end{array}$$

Let $G = \{i \in N : s_i = 0\}$ be the coalition of voters who do not exert effort in equilibrium, which implies $x = \hat{x}^i$ and $\nabla u_i(x) = 0$. For all other voters, the constraint qualification holds at $y = s_i$, so the necessary first order condition is that there exists $\lambda_i \geq 0$ such that

$$\begin{aligned} \nabla u_i(x) &= 2\lambda_i s_i \\ \lambda_i(\eta^2 - \|s_i\|^2) &= 0. \end{aligned}$$

Thus, we have $\|\nabla u_i(x)\| = 2\lambda_i \|s_i\|$, and $\lambda_i = 0$ holds if and only if $x = \hat{x}^i$. Letting $H = \{i \in N \setminus G : \lambda_i > 0\}$, we then have

$$0 = \sum_{i=1}^n s_i = \sum_{i\in H} \frac{1}{2\lambda_i}\nabla u_i(x) + \sum_{i\notin H} s_i. \tag{5}$$

By complementary slackness, it follows that for all $i \in H$, we have $\|s_i\| = \eta$. Thus, for all $i \in H$, we have $\|\nabla u_i(x)\| = 2\eta\lambda_i$, and this implies

$$\sum_{i=1}^n p^i(x) = \sum_{i\in H} p^i(x) = \sum_{i\in H} \frac{1}{2\eta\lambda_i}\nabla u_i(x) = -\frac{1}{\eta}\sum_{i\notin H} s_i,$$

where the last equality follows from (5). Therefore, using the fact that $\|s_i\| \leq \eta$ for all $i \notin H$, we have

$$\left\| \sum_{i=1}^n p^i(x) \right\| = \left\| \frac{1}{\eta}\sum_{i\notin H} s_i \right\| \leq \frac{1}{\eta}\sum_{i\notin H} \|s_i\| \leq \#\{i \in N : x = \hat{x}^i\},$$

and we conclude that x is a directional equilibrium.

For the converse, assume voter utility functions are quasi-concave, let x be a directional equilibrium, and consider any $\eta > 0$ with $D_{n\eta}(x) \subseteq X$. Define the strategy profile s by $s_i = \eta p^i(x)$ for all i, and note that $\sum_{i=1}^n s_i = 0$. To see that s is a Nash equilibrium of the η-contest game, consider any voter i. If $x = \hat{x}^i$, then clearly $s_i = 0 = \eta p^i(x)$ is a best response. Otherwise, if $x \neq \hat{x}^i$, then we return to the above constrained maximization problem,

$$\begin{array}{c} \max_{y \in \mathbb{R}^m} u_i \left(x + (\sum_{j:j \neq i} s_j) + y \right) \\ \text{s.t. } \|y\|^2 \leq \eta^2. \end{array}$$

Note that the constraint function $g(y) = \|y\|^2$ is quasi-convex, with $\nabla g(y) = 2y$. Setting $y = s_i$ and $\lambda_i = \frac{1}{2}\|\nabla u_i(x)\| \geq 0$, we have

$$\begin{array}{rcl} \nabla u_i(x) & = & 2 s_i \lambda_i \\ \lambda_i(\eta^2 - \|s_i\|^2) & = & 0. \end{array}$$

Since u_i is quasi-concave and has non-zero gradient at x, the first order condition is sufficient for a global maximum. Thus, s_i is a best response to the strategies of other voters, and x is non-cooperatively η-stable.

It is straightforward to deduce that s_i specified above is in fact the unique best response for voter i, for suppose there is another best response s_i'. Then in the context of the above constrained maximization problem, there are two optima, $y = s_i$ and $y' = s_i'$. Note that s_i is the unique maximizer of the linear function $f(y) = y \cdot p^i(x)$ subject to $y \in D_\eta(0)$. Since u_i is quasi-concave and $u_i(x + s_i' + \sum_{j \neq i} s_j) \geq u_i(x)$, it follows that $p^i(x) \cdot (s_i' - s_i) \geq 0$, but then s_i' is an additional maximizer of $f(y)$ subject to $y \in D_\eta(0)$, a contradiction. We conclude that x is non-cooperatively strictly η-stable. □

Proof of Theorem 8: Let x be a directional equilibrium, let $\eta > 0$ be such that $D_{n\eta}(x) \subseteq X$, and suppose there is a Nash equilibrium s of the η-contest game at x such that $\sum_j s_j \neq 0$. Let $y = x + \sum_j s_j$, and let $t = \frac{1}{\|\sum_j s_j\|} \sum_j s_j = \frac{1}{\|y - x\|}(y - x)$ be the direction pointing from x to y. Since x is a directional equilibrium, we have $t \cdot \sum_{i=1}^n p^i(x) = 0$. For each voter i, Lemma 1 implies that $t \cdot p^i(y) \leq t \cdot p^i(x)$, and thus we have

$$t \cdot \sum_{i=1}^n p^i(y) \quad \leq \quad 0.$$

In fact, we can strengthen this inequality to account for voters with ideal point at y. Note that for a voter i with $\hat{x}^i = y$, we have,

$$t \cdot p^i(x) \;=\; \frac{1}{\|y-x\|} t \cdot (y-x) \;=\; 1,$$

where we use $t = \frac{1}{\|y-x\|}(y-x)$. Similarly, for a voter i with $\hat{x}^i = x$, we have $t \cdot p^i(y) = -1$. Letting $G = \{i \in N : y = \hat{x}^i\}$ and $H = \{i \in N : x = \hat{x}^i\}$, we can then use Lemma 1 to deduce

$$t \cdot \sum_{i=1}^{n} p^i(y) \;\leq\; t \cdot \sum_{i \in G} p^i(x) - \#G + t \cdot \sum_{i \in H} p^i(x) - \#H + t \cdot \sum_{i \notin G \cup H} p^i(x) \;=\; -\#G,$$

where the last equality uses the assumption that x is a directional equilibrium. But as in the proof of the first part of Theorem 7, for each voter i, s_i solves

$$\begin{array}{c} \max_{y \in \mathbb{R}^m} u_i(x + \sum_{j:j \neq i} s_j + y) \\ \text{s.t. } \|y\|^2 \leq \eta^2, \end{array}$$

and for all $i \notin G$, we thus have $s_i = \eta p^i(y)$. Since $t \cdot p^i(y) = 0 \geq t \cdot s_i - \eta$ for all $i \in G$, we then have

$$t \cdot \sum_{i=1}^{n} p^i(y) \;\geq\; \frac{1}{\eta} t \cdot \sum_{i=1}^{n} s_i - \#G \;=\; \frac{\|\sum_j s_j\|}{\eta} - \#G \;>\; -\#G,$$

a contradiction. □

Proof of Theorem 9: Assume x is a non-degenerate directional equilibrium, and to simplify accounting, assume $D_n(x) \subseteq X$. We initially construct a strategy profile s for the 1-contest game at x as follows. For all $i \in N$ such that $\nabla u_i(x) \neq 0$, we set $s_i^* = p^i(x)$. Letting $G = \{i \in N : \nabla u_i(x) \neq 0\}$, the definition of directional equilibrium implies $\| \sum_{i \in G} p^i(x) \| \leq \#\{i \in N : x = \hat{x}^i\}$, so we can specify s_i^* for all $i \notin G$ such that $\|s_i^*\| \leq 1$ and $\sum_{i \notin G} s_i^* = -\sum_{i \in G} p^i(x)$, which implies $\sum_i s_i^* = 0$. Consider the following maximization problem.

$$\begin{array}{c} \max_{y \in \mathbb{R}^m} u_i\left(x + (\sum_{j:j \neq i} s_j^*) + y\right) \\ \text{s.t. } \|y\|^2 \leq 1. \end{array}$$

For all $i \in G$, set $\lambda_i = \frac{1}{2}\|\nabla u_i(x)\|$, and note that the first order condition at s_i^*,

$$\begin{aligned} \nabla u_i(x) &= 2\lambda_i s_i^* \\ \lambda_i(1 - \|s_i^*\|^2) &= 0, \end{aligned}$$

is satisfied. Furthermore, because the constraint function is quadratic and x is non-degenerate, the second order sufficient condition holds, so s_i^* is a strict local solution to the above problem. Thus, there exists $\epsilon_i' > 0$ such that for all s_i' with $\|s_i'\| \leq 1$ and $\|s_i' - s_i^*\| < \epsilon_i'$, we have

$$u_i(x) \;\; > \;\; u_i\left(x + (\sum_{j:j\neq i} s_j^*) + s_i'\right).$$

Set $\epsilon_i = \epsilon_i'/3$. For all $i \notin G$, we have $x = \hat{x}^i$. Of course, if $G = \emptyset$, then x is clearly non-cooperatively strictly ϵ-stable for all $\epsilon > 0$.

Assuming, then, that $G \neq \emptyset$, we will show that x is non-cooperatively strictly ϵ-stable for $\epsilon < \eta = \min\{\epsilon_i : i \in G\}$. We specify the strategy profile s in the ϵ-contest game such that $s_i = \epsilon s_i^*$ for all voters i. To see that s_i is the unique best response for each $i \in G$, consider any $\tilde{s}_i$ such that $\|\tilde{s}_i\| \leq \epsilon$. Then define

$$s_i' \;\; = \;\; \tilde{s}_i - (1-\epsilon)\sum_{j:j\neq i} s_j^*,$$

and note that since $s_i^* = -\sum_{j:j\neq i} s_j^*$, we have

$$\|s_i'\| \;\; = \;\; \left\|\tilde{s}_i - (1-\epsilon)\sum_{j:j\neq i} s_j^*\right\| \;\; \leq \;\; \|\tilde{s}_i\| + (1-\epsilon)\left\|\sum_{j:j\neq i} s_j^*\right\| \;\; \leq \;\; 1,$$

and furthermore,

$$\|s_i' - s_i^*\| \;\; = \;\; \left\|\tilde{s}_i - (1-\epsilon)\sum_{j:j\neq i} s_j^* - s_i^*\right\| \;\; = \;\; \|\tilde{s}_i + \epsilon s_i^*\| \;\; \leq \;\; 2\epsilon \;\; < \;\; \epsilon_i'.$$

By construction of ϵ_i', we therefore have

$$\begin{aligned} u_i(x) \;\; &> \;\; u_i\left(x + (\sum_{j:j\neq i} s_j^*) + s_i'\right) \;\; = \;\; u_i\left(x + \epsilon(\sum_{j:j\neq i} s_j^*) + \tilde{s}_i\right) \\ &= \;\; u_i\left(x + (\sum_{j:j\neq i} s_j) + \tilde{s}_i\right), \end{aligned}$$

and we conclude that s is a strict Nash equilibrium of the ϵ-contest game at x with $\sum_i s_i = 0$, as required. □

Example of unstable directional equilibrium: Suppose toward a contradiction that there exists $\epsilon > 0$ such that there is a Nash equilibrium of the ϵ-local game at

$(0,0)$ such that $\sum_{i=1}^{3} s_i = 0$. In equilibrium, each voter i solves the following maximization problem.

$$\begin{array}{c}\max_{(x,y)\in\mathbb{R}^2} u_i\Big((\sum_{j:j\neq i} s_j^*) + (x,y)\Big)\\ \text{s.t. } x^2+y^2 \leq \epsilon^2.\end{array}$$

Clearly, the solution for voters 1 and 3 is non-zero, so the constraint qualification holds for both. For voter 1, for example, the necessary first order condition is that there exists $\lambda_1 \geq 0$ satisfying

$$\begin{aligned}\nabla u_1(0,0) &= 2\lambda_1 s_1^*\\ \lambda_1(\epsilon^2 - \|s_1^*\|^2) &= 0.\end{aligned}$$

Thus, we obtain that s_1^* points in the direction of $(0,1)$ and has norm ϵ, i.e., $s_1^* = (0,\epsilon)$. Similarly, $s_3^* = (0,-\epsilon)$, and thus $s_2^* = (0,0)$.

Note, however, that while s_1^* solves the necessary first order condition for voter 1, it is not a best response to $s_2^* + s_3^* = (0,-\epsilon)$. Using the functional form of the voter's utility and changing variables, voter 1 solves

$$\begin{array}{c}\max_{(x,y)\in\mathbb{R}^2} y + x^3 + x^3\sin(1/x)\\ \text{s.t. } \|(x,y)-(0,-\epsilon)\| \leq \epsilon,\end{array}$$

and s_1^* corresponds to a potential solution $(0,\epsilon)$. Thus, the outcome from following s_1^* is the origin, and voter 1's utility is zero. But consider the voter's utility from (x',y') with $x' = \alpha > 0$ and $y' = -\epsilon + \sqrt{\epsilon^2 - \alpha^2}$. This is the quantity

$$-\epsilon + \sqrt{\epsilon^2-\alpha^2} + \alpha^3 + \alpha^3\sin(1/\alpha),$$

which is positive if

$$1 \quad > \quad \frac{\epsilon - \sqrt{\epsilon^2-\alpha^2}}{\alpha^3 + \alpha^3\sin(1/\alpha)}.$$

Applying L'Hôpital's rule twice, the limit of the right-hand side of the above inequality for small $\alpha > 0$ is the limit of

$$\frac{(\epsilon^2-\alpha^2)^{-\frac{1}{2}} + \alpha^2(\epsilon^2-\alpha^2)^{-\frac{3}{2}}}{6\alpha + 6\alpha\sin(1/\alpha) - 2\cos(1/\alpha) + \frac{1}{\alpha}\sin(1/\alpha)}$$

as $\alpha \to 0$. Choosing the sequence $\{\alpha_k\}$ with $\alpha_k = 2/\pi k$, the limit of the right-hand side is therefore zero. We conclude that for any given $\epsilon > 0$, voter 1's utility from choosing (x',y') with $\alpha > 0$ appropriately small is positive, contradicting the assumption that s_1^* is a best response.

References

[1] J. Banks, J. Duggan, and M. Le Breton (2006) "Social Choice and Electoral Competition in the General Spatial Model," *Journal of Economic Theory,* 126: 194–234.

[2] N. Baranchuk and P. Dybvig (2009) "Consensus in Diverse Corporate Boards," *Review of Financial Studies,* 22: 715–747.

[3] D. Benjamin, O. Heffetz, M. Kimball, and N. Szembrot (2013) "Aggregating Local Preferences to Guide Marginal Policy Adjustments," *American Economic Review: Papers & Proceedings,* 103: 605–610.

[4] D. Benjamin, G. Carroll, O. Heffetz, and M. Kimball (2014) "Aggregating Local Preferences to Guide Policy," unpublished paper.

[5] R. Brady and C. Chambers (2015) "Spatial Implementation," *Games and Economic Behavior,* 94: 200–205.

[6] R. Brady and C. Chambers (2016) "A Spatial Analogue of May's Theorem," *Social Choice and Welfare,* 47: 127–139.

[7] B. Brown (1983) "Statistical Uses of the Spatial Median," *Applied Statistics,* 23: 325–328.

[8] D. Cervone, R. Dai, D. Gnoutcheff, G. Lanterman, A. Mackenzie, A. Morse, N. Srivastava, W. Zwicker (2012) "Voting with Rubber Bands, Weights, and Strings," *Mathematical Social Sciences,* 64: 11–27.

[9] H. Chung and J. Duggan (2017) "Directional Equilibrium," unpublished paper.

[10] J. Duggan (2013) "Uncovered Sets," *Social Choice and Welfare,* 41: 489–535.

[11] C. Gini and L. Galvani (1929) "Di Talune Estensioni dei Concetti di Media ai Caratteri Qualitativi," *Metron,* 8: 3–209.

[12] J. Gower (1974) "The Mediancentre," *Journal of the Royal Statistical Society,* 23: 466-470.

[13] B. Grofman, G. Owen, N. Noviello, and A. Glazer (1987) "Stability and Centrality of Legislative Choice in the Spatial Context," *American Political Science Review,* 81: 539–553.

[14] J. Haldane (1948) "Note on the Median of a Multivariate Distribution," *Biometrika,* 35: 414–415.

[15] A. Hylland and R. Zeckhauser (1980) "A Mechanism for Selecting Public Goods when Preferences Must Be Elicited," unpublished paper.

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

[17] G. Owen and L. Shapley (1989) "Optimal Location of Candidates in Ideological Space," *International Journal of Game Theory,* 18: 339–356.

[18] J. Patty (2007) "Generic Difference of Expected Vote Share and Probability of Victory Maximization in Simple Plurality Elections with Probabilistic Voters," *Social Choice and Welfare,* 29: 149–173.

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

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

[21] K. Shepsle (1979) "Institutional Arrangements and Equilibrium in Multidimensional Voting Models," *American Journal of Political Science,* 23: 27–59.

[22] C. Small (1990) "A Survey of Multidimensional Medians," *International Statistical Review,* 58: 263–277.

[23] A. Weber (1909) "Über den Standort der Industrien," Tubingen. English translation by C.J. Friedrich (1929) *Alfred Webver's Theory of Location of Industries,* Chicago: University of Chicago Press.

[24] E. Van Wesep (2010) "Defensive Politics," unpublished paper.

[25] A. Wuffle, S. Feld, G. Owen, and B. Grofman (1989) "Finagle's Law and the Finagle Point, a New Solution Concept for Two-candidate Competition in Spatial Voting Games without a Core," *American Journal of Political Science,* 33: 348–375.