# Computing Pure-Strategy Nash Equilibria in a Two-Party Policy Competition: Existence and Algorithmic Approaches

Extended Abstract

Chuang-Chieh Lin
Department of Computer Science and Engineering,
National Taiwan Ocean University
Keelung City, Taiwan
josephcclin@mail.ntou.edu.tw

Chi-Jen Lu
Institute of Information Science,
Academia Sinica
Taipei City, Taiwan
cjlu@iis.sinica.edu.tw

Po-An Chen
Institute of Information Management,
National Yang Ming Chiao Tung University
Hsinchu City, Taiwan
poanchen@nycu.edu.tw

Chih-Chieh Hung
Department of Management Information Systems,
National Chung Hsing University
Taichung City, Taiwan
smalloshin@nchu.edu.tw

## ABSTRACT

We study two-party policy competition as a two-player game in which each party selects a policy vector from a compact subset of $\mathbb{R}^k$ and voters evaluate policies via inner products with their preference vectors. We model electoral uncertainty by an affine isotone winning-probability function of the total utility difference across all voters, and define payoffs as supporters' expected utility. We prove existence of a pure-strategy Nash equilibrium (PSNE) in both one- and multi-dimensional settings, with a closed-form characterization in one dimension. Although the game is not monotone in general, experiments (see the full version) suggest decentralized gradient-based dynamics typically converge quickly to approximate PSNE. Finally, we give a polynomial-time grid-based algorithm to compute an $\varepsilon$-approximate PSNE.

## KEYWORDS

Pure-Strategy Nash equilibrium; Continuous Games; Policy Competition; Fixed-Point; Gradient Dynamics.

**ACM Reference Format:**

Chuang-Chieh Lin, Chi-Jen Lu, Po-An Chen, and Chih-Chieh Hung. 2026. Computing Pure-Strategy Nash Equilibria in a Two-Party Policy Competition: Existence and Algorithmic Approaches: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/RUHW5194

## 1 INTRODUCTION AND MOTIVATION

Political competition has long been studied through the *Spatial Theory of Voting* [3, 6, 9, 12, 15], and Duverger's law suggests that *plurality voting* promotes a two-party system [2, 5]. Motivated by these works, Lin et al. [7] modeled two-party competition as a two-player game in which each party nominates a candidate from a finite set, with winning probability proportional to total voter utility and payoff defined as supporters' expected utility, establishing PSNE existence and PoA bounds.

In this paper, we extend Lin et al.'s work [7] from discrete candidates to continuous policy proposals. Each party chooses a policy vector in a compact subset of $\mathbb{R}^k$, and a voter's utility is the inner product between the policy and the voter's preference vector, so a policy may benefit supporters of either party. This continuous-action model yields an infinite strategy space, where PSNE existence is nontrivial. Though mixed equilibria are guaranteed only for finite games [10, 11], PSNE may fail to exist even for two players (e.g., Matching Pennies), and deciding PSNE existence in an $n$-player finite game is NP-complete [1].

### 1.1 Our Contributions

Our model complements classic distance-based spatial voting. Using inner products, utility captures both *directional alignment* (beneficial vs. harmful) and *intensity*, whereas distance-based models penalize magnitude without distinguishing directions. For example, let a voter have $\mathbf{q} = (0.1, -0.1)$ and consider $\mathbf{z}_B = (0.3, -0.3)$ and $\mathbf{z}_H = (-0.1, 0.1)$. Although $\|\mathbf{z}_B - \mathbf{q}\| = \|\mathbf{z}_H - \mathbf{q}\| = \sqrt{0.08}$, we have $\mathbf{q}^\top \mathbf{z}_B = 0.06 > 0$ but $\mathbf{q}^\top \mathbf{z}_H = -0.02 < 0$, showing that inner products distinguish aligned and misaligned policies that distance cannot. This formulation also highlights a natural trade-off between satisfying supporters and compromising to attract all voters. Moreover, unlike high-dimensional majority-rule spatial models that can be unstable [4], our isotonicity-based formulation guarantees PSNE existence and yields a computable equilibrium benchmark. Our contributions are as follows.

- We give a closed-form PSNE characterization in the one-dimensional setting.
- We prove PSNE existence in the multi-dimensional setting.
- We show the game is not monotone in general; nevertheless, simulations suggest vanilla gradient ascent often converges to approximate PSNE in practice.
- We design a polynomial-time grid-based algorithm to compute an $\varepsilon$-approximate PSNE.



*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus.* © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/RUHW5194

For an extended version of this work with further details refer to [8].

## 2 THE MODEL AND PRELIMINARIES

We model two-party policy competition as a two-player game between parties $A$ and $B$. Each party $X \in \{A, B\}$ chooses a policy vector $\mathbf{z}_X \in S \subseteq \mathbb{R}^k$, where $S := \{\mathbf{z} \in [-1,1]^k : \|\mathbf{z}\| \leq 1\}$, and we write $\mathbf{z} = (\mathbf{z}_A, \mathbf{z}_B)$. Voters are partitioned into supporters $V_A$ and $V_B$ with $V = V_A \uplus V_B$. Each voter $v$ has a preference vector $\mathbf{q}_v \in S$; define $Q_X := \sum_{v \in V_X} \mathbf{q}_v$ and $Q := Q_A + Q_B$ (normalized so $\|Q_A\|, \|Q_B\| \leq 1$). The utility that policy $\mathbf{z}$ provides to supporters of $X$ is $u_X(\mathbf{z}) := \langle \mathbf{z}, Q_X \rangle = \mathbf{z}^\top Q_X$.

Following the isotonicity hypothesis of [7], we model the winning probability of $A$ by the affine isotonic function

$$p_A(\mathbf{z}) = \frac{1}{2} + \frac{1}{8}(\mathbf{z}_A - \mathbf{z}_B)^\top Q, \qquad p_B(\mathbf{z}) = 1 - p_A(\mathbf{z}). \qquad (1)$$

The payoff of party $A$ is the expected utility received by its supporters, that is, $R_A(\mathbf{z}) = p_A(\mathbf{z})\, \mathbf{z}_A^\top Q_A + (1 - p_A(\mathbf{z}))\, \mathbf{z}_B^\top Q_A$, and $R_B(\mathbf{z})$ is defined symmetrically. Each party seeks to maximize its own payoff by choosing $\mathbf{z}_X \in S$.

## 3 THE EXISTENCE OF PSNE

### 3.1 The One-Dimensional Setting

For $k = 1$, policies are scalars $z_A, z_B \in [-1, 1]$. Under (1) and the payoff definition, party $A$'s payoff can be written as

$$R_A(z_A, z_B) = \tfrac{1}{2}(z_A + z_B)Q_A + \tfrac{1}{8}\, QQ_A\, (z_A - z_B)^2,$$

and $R_B$ is defined symmetrically.

Theorem 1. *The one-dimensional two-party policy competition game admits a PSNE. Moreover, a PSNE can be expressed in closed form (see the extended version [8]).*

### 3.2 The Multi-Dimensional Setting

We now consider $k \geq 1$. Since payoffs depend on inner products with $Q_A$ and $Q_B$, only the projection of each policy onto $\operatorname{span}\{Q_A, Q_B\}$ matters; hence we may restrict to this two-dimensional subspace. In this plane, we represent each policy by polar coordinates $(r_X, \theta_X)$, where $r_X = \|\mathbf{z}_X\| \in [0, 1]$ and $\theta_X$ is the angle relative to the corresponding supporter vector.

Lemma 1. *In every best response (and thus at any PSNE), both parties choose maximal strength: $r_A = r_B = 1$.*

Theorem 2. *For every $k \geq 1$, the two-party policy competition game admits a PSNE. The proof reduces the game to angular best responses in* $\operatorname{span}\{Q_A, Q_B\}$ *and applies fixed-point arguments. Details of the proof appear in the extended version [8].*

## 4 BAD AND GOOD NEWS FOR GRADIENT-BASED DYNAMICS

*Pseudo-gradient view.* For a differentiable game, define the pseudo-gradient map $F(\mathbf{x}) := (\nabla_{x_i} u_i(x_i, \mathbf{x}_{-i}))_{i=1}^n$ [14]. A profile $\mathbf{x}^*$ is a PSNE if and only if it solves the variational inequality $F(\mathbf{x}^*)^\top(\mathbf{y} - \mathbf{x}^*) \leq 0$ for all $\mathbf{y}$ in the feasible set.

Although our game admits PSNE, its pseudo-gradient map needs not be monotone. In particular, even in the reduced two-dimensional angular representation (with $r_A = r_B = 1$ and variables $x = \cos\theta_A$, $y = \cos\theta_B$), we construct $(x_1, y_1)$ and $(x_2, y_2)$ such that $\big(F(x_1, y_1) - F(x_2, y_2)\big)^\top \big((x_1, y_1) - (x_2, y_2)\big) > 0$, which disproves monotonicity (and thus rules out any $\lambda$-cocoercivity). Consequently, standard convergence guarantees for gradient methods do not apply. (See the extended version for details and numerical values.)

Despite the above, a simple decentralized projected gradient ascent performs well empirically: $\mathbf{z}_X^{(t)} \leftarrow \Pi_S\big(\mathbf{z}_X^{(t-1)} + \eta_t \nabla_{\mathbf{z}_X} R_X(\mathbf{z}^{(t-1)})\big)$, for $X \in \{A, B\}$, with a decaying step size (e.g., $\eta_t = t^{-0.75}$). Across random preference instances and random initial policies, the dynamics typically converge within a few thousand iterations, and the terminal profiles are approximate PSNE in the vast majority of runs (see the extended version for experimental details). We also tested an extragradient variant, but did not observe consistent improvements.

## 5 DISCRETIZATION COMES TO THE RESCUE

Since the strategy space is compact, we can compute approximate equilibria via discretization. Working in the reduced angular representation $(\theta_A, \theta_B) \in [0, \rho_A] \times [0, \rho_B]$ (with $r_A = r_B = 1$), the pseudo-gradient map $F(\theta_A, \theta_B)$ is Lipschitz with constant $\mathcal{L} = O(\|Q_A\| + \|Q_B\|)$ on this domain. Therefore, a uniform grid of spacing $h = \Theta(\varepsilon/\mathcal{L})$ suffices: any unilateral deviation can be approximated by a nearby grid point with payoff loss at most $\mathcal{L}h$.

This yields a simple grid-based algorithm (`GBA-PSNE`): discretize each party's feasible angles, compute approximate best responses on the grid, and return a profile with deviation gain at most $\varepsilon$ (details in the full version).

Theorem 3. `GBA-PSNE` *outputs an $\varepsilon$-PSNE in time polynomial in the input size and $1/\varepsilon$. Moreover, using unimodality of one-dimensional best-response objectives on the grid, the search can be implemented in $O\big(nk + \frac{1}{\varepsilon}\log\frac{1}{\varepsilon}\big)$ time.*

## 6 CONCLUDING REMARKS

We model two-party policy competition as a continuous-action game with inner-product utilities and an isotone winning-probability rule. We validate isotonicity empirically and prove PSNE existence in both one- and multi-dimensional settings, complementing classic spatial models that may lack equilibrium in higher dimensions [13, 16]. Although the game is not monotone in general, projected gradient ascent performs well empirically, and a grid-based algorithm yields guaranteed $\varepsilon$-approximate PSNE in polynomial time.

*Future directions.* Alternative isotone winning-probability functions; multi-party competition; learning and strategic behavior under partial information.

## ACKNOWLEDGMENTS

This work is funded and supported by the National Science and Technology Council, Taiwan, under grant nos. NSTC 112-2221-E-032-018-MY3, NSTC 113-2628-HA49-002-MY3, and NSTC 114-2221-E005-058. This work is also partially supported by NSTC 114-2634-F-005-002- through the Smart Sustainable New Agriculture Research Center (SMARTer).

# REFERENCES


[1] C. Àlvarez, J. Gabarro, and M. Serna. 2011. Equilibria problems on games: Complexity versus succinctness. *J. Comput. System Sci.* 77, 6 (2011), 1172–1197. https://doi.org/10.1016/j.jcss.2011.01.001
[2] A. Dellis. 2013. The two-party system under alternative voting procedures. *Social Choice and Welfare* 40 (2013), 263–284.
[3] A. Downs. 1957. *An Economic Theory of Democracy.* Harper and Row, New York.
[4] J. Duggan. 2008. Candidate objectives and electoral equilibrium. In *The Oxford handbook of political economy*, B. Weingast and D. Wittman (Eds.). Oxford University Press.
[5] M. Duverger. 1954. *Political Parties: their organization and activity in the modern state.* Methuen & Co.; John Wiley & Sons London, New York.
[6] H. Hotelling. 1929. Stability in competition. *The Economic Journal* 39, 153 (1929), 41–57.
[7] C.-C. Lin, C.-J. Lu, and P.-A. Chen. 2021. How good is a two-party election game? *Theoretical Computer Science* 871 (2021), 79–93. https://doi.org/10.1016/j.tcs.2021.04.013
[8] C.-C. Lin, C.-J. Lu, P.-A. Chen, and C.-C. Hung. 2025. Computing Pure-Strategy Nash Equilibria in a Two-Party Policy Competition: Existence and Algorithmic Approaches. arXiv:2512.22552 [cs.GT] https://arxiv.org/abs/2512.22552
[9] M. D. López, J. Rodrigo, and I. Lillo. 2007. Two-party political competition: a geometric study of the Nash equilibrium in a weighted case. *Applied Mathematical Sciences* 1, 55 (2007), 2715–2725.
[10] J. F. Nash. 1950. Equilibrium points in $n$-person games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49. https://doi.org/10.1073/pnas.36.1.48
[11] J. F. Nash. 1951. Noncooperative games. *Annals of Mathematics* 54 (1951), 289–295. https://doi.org/10.2307/1969529
[12] T. R. Palfrey. 1984. Spatial equilibrium with entry. *The Review of Economic Studies* 51, 1 (1984), 139–156.
[13] C. R. Plott. 1967. A notion of equilibrium and its possibility under majority rule. *The American Economic Review* 57, 4 (1967), 787–806.
[14] J. B. Rosen. 1965. Existence and uniqueness of equilibrium points for concave $N$-person games. *Econometrica* 33, 3 (1965), 520–534. https://doi.org/10.2307/1911749
[15] S. Weber. 1992. On hierarchical spatial competition. *Review of Economic Studies* 59, 2 (1992), 407–425.
[16] D. Xefteris. 2017. Multidimensional electoral competition between differentiated candidates. *Games and Economic Behavior* 105 (2017), 112–121. https://doi.org/10.1016/j.geb.2017.07.005