Computing PSNE in a Two-Party Policy Competition: Existence and Algorithmic Approaches

Chuang-Chieh Lin, Chi-Jen Lu, Po-An Chen, Chih-Chieh Hung

Department of Computer Science and Engineering, National Taiwan Ocean University, TW
Institute of Information Science, Academia Sinica, TW
Institute of Information Management, National Yang Ming Chiao Tung University, TW
Department of Management Information Systems, National Chung Hsing University, TW
josephcclin@mail.ntou.edu.tw;cjlu@iis.sinica.edu.tw;poanchen@nycu.edu.tw;
smalloshin@nchu.edu.tw

Abstract

We formulate two-party policy competition as a two-player continuous-action game. Each party selects a policy vector in a compact domain; voters evaluate policies via inner products with preference vectors. Election outcome is uncertain and is modeled by an isotone winning-probability function of the aggregate utility advantage. A party's payoff is the expected utility received by its supporters. We prove existence of a pure-strategy Nash equilibrium (PSNE) in both one- and multi-dimensional settings (with closed forms in special cases), show the induced pseudo-gradient operator is not monotone in general, empirically study decentralized gradient-based learning, and give a grid-based algorithm that computes an ε-approximate PSNE in time polynomial in the input size and $1/\varepsilon$.

Motivation

- Classic spatial voting models are often distance-based and can lack stability in higher dimensions.
- We study a *macro* game: parties are strategic players choosing multi-issue platforms.
- Inner-product utilities capture *directional alignment* (beneficial vs. harmful) and *intensity*.
- Isotonicity hypothesis: higher aggregate utility advantage implies higher winning probability.

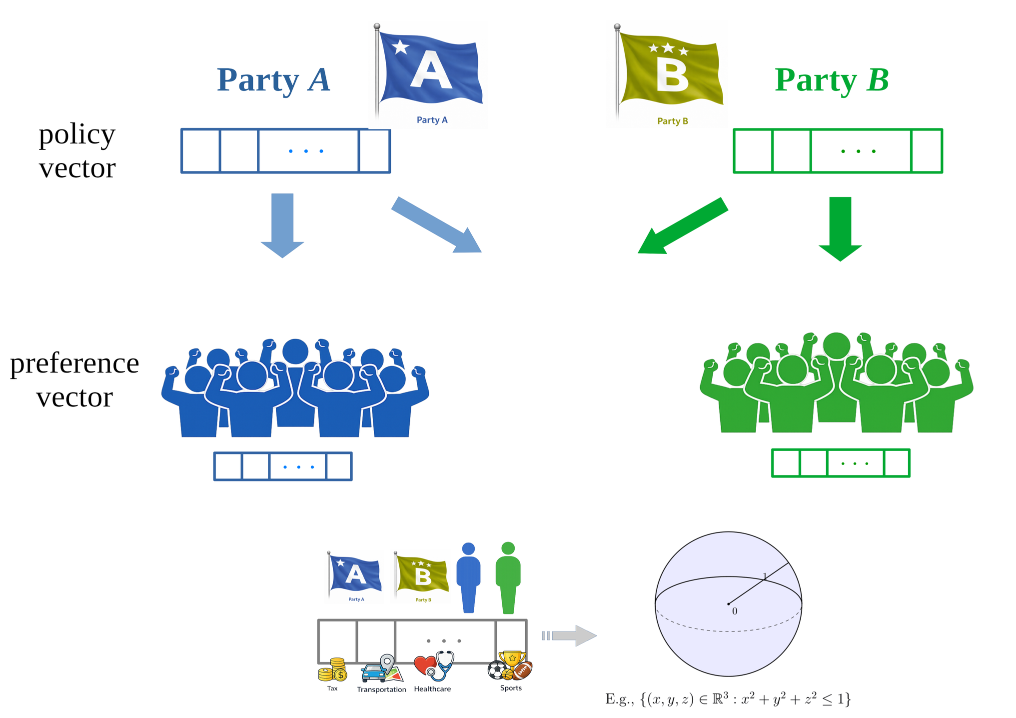

Figure 1: *Two-party policy competition: policy vectors vs. voter preferences.*

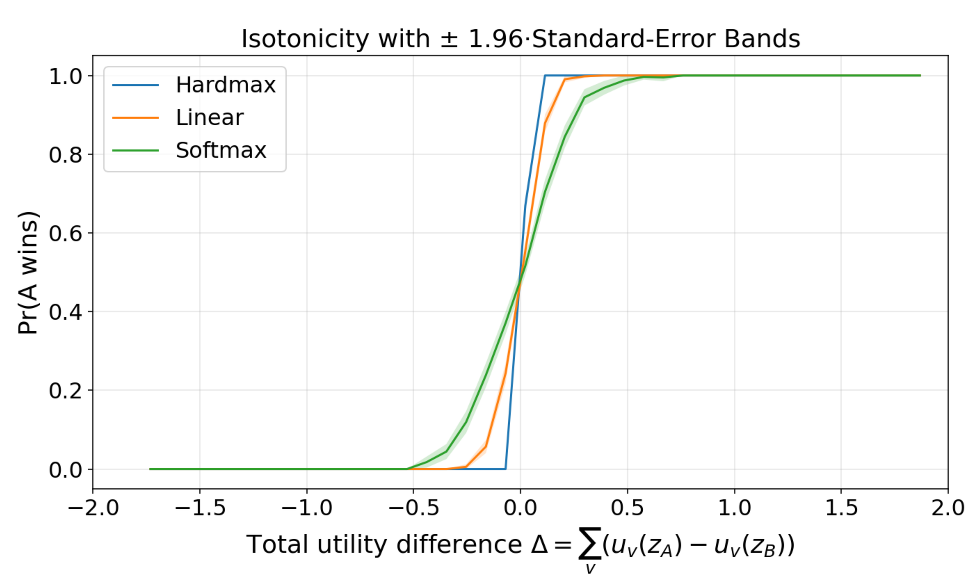

Figure 2: *Utility difference vs. probability that party A wins. Illustration of the isotonicity.*

The Model & Assumptions

- Policy domain: $S := \{\mathbf{z} \in [-1,1]^k : \|\mathbf{z}\| \le 1\}$. Parties choose $\mathbf{z}_A, \mathbf{z}_B \in S$.
- Supporters V_A, V_B with preference vectors $\mathbf{q}_v \in S$. Define $Q_A = \sum_{v \in V_A} \mathbf{q}_v$, $Q_B = \sum_{v \in V_B} \mathbf{q}_v$, $Q = Q_A + Q_B$.
- Winning probability (linear link):

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

  More generally, $p_A(\mathbf{z}) = \mu((\mathbf{z}_A - \mathbf{z}_B)^\top Q)$ for an isotone link μ.
- Payoff: expected supporter utility

$$R_A(\mathbf{z}) = p_A(\mathbf{z})\,\mathbf{z}_A^\top Q_A + (1 - p_A(\mathbf{z}))\,\mathbf{z}_B^\top Q_A \quad \text{(and sym. for } B\text{).}$$

Main Results (I): PSNE Existence Guarantee

- **PSNE existence:** Under mild conditions on the isotone link (full version), the game admits a PSNE (via Kakutani's fixed-point theorem).
- **1D / collinear case:** Closed-form PSNE characterization.
- **2D reduction:** Only the projection to $\mathrm{span}\{Q_A, Q_B\}$ matters. Use polar coordinates (r_X, θ_X).
- **Maximal strength:** In best responses (and thus at PSNE), $r_A = r_B = 1$.

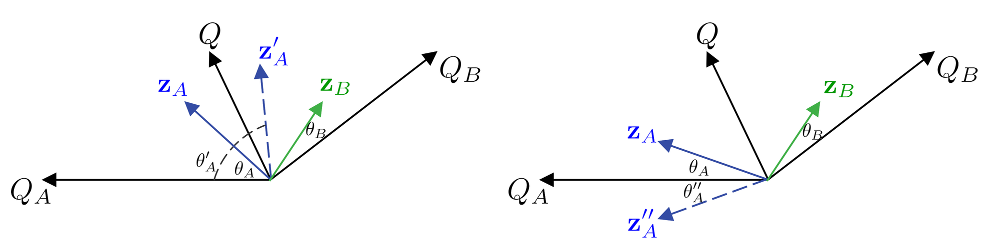

Figure 3: *Angle restriction: it suffices for party A to consider policies in a wedge between Q_A and Q.*

Main Results (II): Decentralized Projected Gradient Ascent

- Pseudo-gradient operator is *not monotone* in general (counterexamples), so standard VI convergence theory does not apply.
- Empirically, decentralized projected gradient ascent often converges quickly and typically ends at an approximate PSNE.

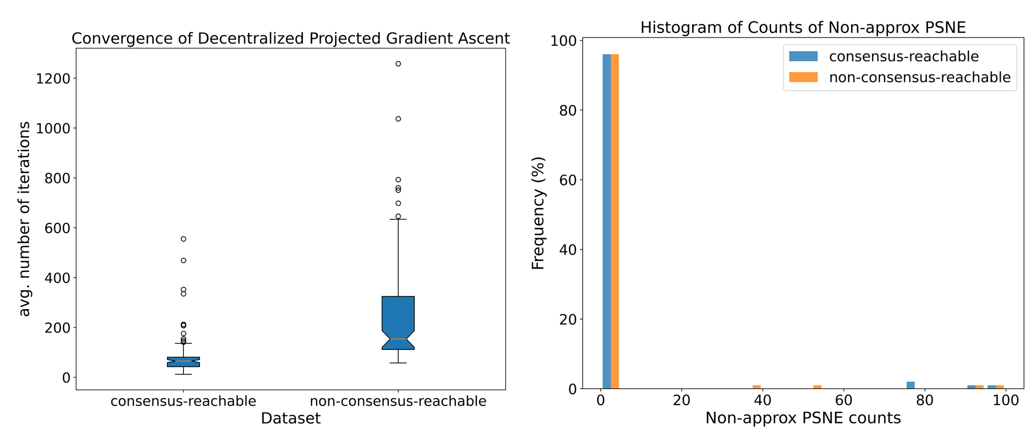

Figure 4: *Left: convergence speed. Right: frequency of convergences that are not approx. PSNE.*

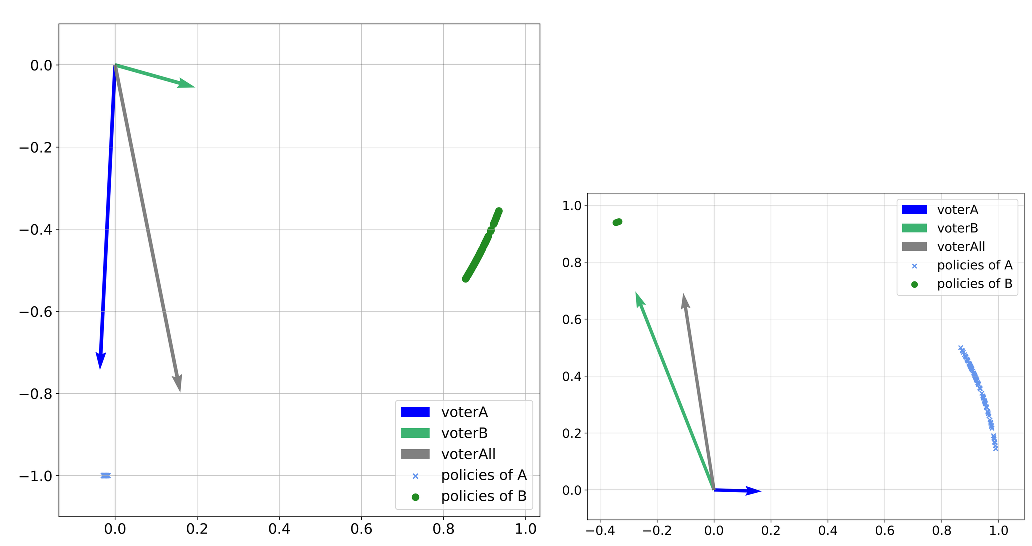

Figure 5: *Decentralized projected gradient ascent may lead to multiple convergence.*

Main Results (III): Grid-Based ε-PSNE (GBA-PSNE)

- Work in reduced angles $(\theta_A, \theta_B) \in [0, \rho_A] \times [0, \rho_B]$ with $r_A = r_B = 1$.
- Lipschitz bound L implies grid spacing $h = \Theta(\varepsilon/L)$ suffices.
- Linear link: exploit unimodality of best-response slices to get $O((1/\varepsilon)\log(1/\varepsilon))$ time search.

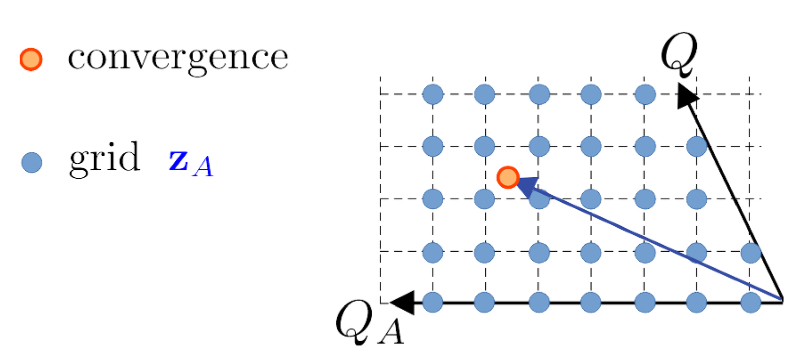

Figure 6: *Discretization and deviation checking on a grid.*

Takeaways & Future Directions

- A continuous-action two-party policy game with inner-product utilities and isotone winning probabilities admits PSNE.
- Discretization yields *provable* ε-PSNE algorithms; gradient learning works well empirically despite non-monotonicity.
- Future work: richer isotone links, multi-party competition, partial information / learning voters.

References

1. H. Hotelling, Stability in competition, *The Economic Journal* 39 (153) (1929) 41–57. doi:10.2307/2224214.
2. A. Downs, An Economic Theory of Democracy, Harper and Row, New York, 1957.
3. M. D. López, J. Rodrigo, I. Lillo, Two-party political competition: a geometric study of the Nash equilibrium in a weighted case, *Applied Mathematical Sciences* 1 (55) (2007) 2715–2725.
4. T. R. Palfrey, Spatial equilibrium with entry, *The Review of Economic Studies* 51 (1) (1984) 139–156. doi:10.2307/2297710.
5. S. Weber, On hierarchical spatial competition, *The Review of Economic Studies* 59 (2) (1992) 407–425. doi:10.2307/2297961.
6. M. Duverger, Political Parties: their organization and activity in the modern state, Methuen & Co.; John Wiley & Sons London, New York, 1954.
7. A. Dellis, The two-party system under alternative voting procedures, *Social Choice and Welfare* 40 (2013) 263–284. doi:10.1007/s00355-011-0597-3.
8. C.-C. Lin, C.-J. Lu, P.-A. Chen, How good is a two-party election game? *Theoretical Computer Science* 871 (2021) 79–93. doi:10.1016/j.tcs.2021.04.013.
9. J. F. Nash, Equilibrium points in n-person games, Proceedings of the National Academy of Sciences 36 (1) (1950) 48–49. doi:10.1073/pnas.36.1.48.
10. J. F. Nash, Noncooperative games, *Annals of Mathematics* 54 (1951) 286–295. doi:10.2307/1969529.
11. C. Àlvarez, J. Gabarro, M. Serna, Equilibria problems on games: Complexity versus succinctness, *Journal of Computer and System Sciences* 77 (6) (2011) 1172–1197. doi:10.1016/j.jcss.2011.01.001.
12. J. Duggan, Candidate objectives and electoral equilibrium, in: B. Weingast, D. Wittman (Eds.), The Oxford Handbook of Political Economy, Oxford University Press, 2008.
13. G. Farina, I. Anagnostides, H. Luo, C.-W. Lee, C. Kroer, T. Sandholm, Near-optimal no-regret learning dynamics for general convex games, in: Advances in Neural Information Processing Systems (NeurIPS'22), 2022.
14. Y. Cai, W. Zheng, Doubly optimal no-regret learning in monotone games, in: International Conference on Machine Learning (ICML'23), 2023, pp. 3507–3524.
15. S. Jelassi, C. Domingo-Enrich, D. Scieur, A. Mensch, J. Bruna, Extragradient with player sampling for faster Nash equilibrium finding, in: Proceedings of the 37th International Conference on Machine Learning (ICML'20), 2020, pp. 4736–4745.
16. E. Gorbunov, N. Loizou, G. Gidel, Extragradient method: $O(1/k)$ last-iterate convergence for monotone variational inequalities and connections with cocoercivity, in: Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS'22), Vol. 151 of Proceedings of Machine Learning Research, 2022, pp. 366–402.

AAMAS Extended Abstract Poster, Cyprus, 25–29 May 2026. Full version: https://arxiv.org/abs/2512.22552