Real Preferences under Arbitrary Norms

Extended Abstract

Joshua Zeitlin
University of Cambridge
United Kingdom
jdz28@cam.ac.uk

Corinna Coupette
Aalto University
Finland
corinna.coupette@aalto.fi

ABSTRACT

Whether the goal is to reach decisions among multiple agents, ensure that AI systems are aligned with human preferences, or design better recommender systems, the problem of translating between (ordinal) rankings and (numerical) utilities arises naturally in many contexts. This task is commonly approached by computing *embeddings*, which represent both the agents doing the ranking (*voters*) and the items to be ranked (*alternatives*) in a shared metric space. Here, ordinal preferences are translated into relationships between pairwise distances. Prior work has established that any collection of rankings with n voters and m alternatives (*preference profile*) can be embedded into d-dimensional Euclidean space for $d \geq \min\{n, m-1\}$ under the Euclidean norm and the Manhattan norm. We show that this holds for *all p-norms* and establish that any *pair* of rankings can be embedded into $\mathbb{R}^2$ under *arbitrary norms*, significantly expanding the reach of spatial preference models.

KEYWORDS

Spatial Preferences, Social Choice Theory, Voting

ACM Reference Format:
Joshua Zeitlin and Corinna Coupette. 2026. Real Preferences under Arbitrary Norms: 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/JWXO7778

1 INTRODUCTION

From selecting winners in real-world elections to reaching decisions in multiagent systems, designing powerful recommendation engines, and improving language models with human feedback: The problem of evaluating alternatives based on preference information collected from multiple parties lies at the heart of diverse practical applications. In many settings, the preference information available is *ordinal*—e.g., because rankings and pairwise comparisons can be intuitively understood by humans and communicated efficiently between agents. At the same time, *metric* representations of preferences tend to allow more expressive evaluations and more efficient computations. As a result, identifying "good" mappings from ordinal to metric preferences constitutes a key practical challenge.

In economics and political science, this challenge is commonly approached via the widely popular *theory of spatial preferences*, which represents agents (e.g., voters) and alternatives (e.g., candidates or policy outcomes) as points in Euclidean space [19, 40]. Here, the underlying norm is often assumed to be Euclidean as well; individual utilities and collective social costs are treated as functions of Euclidean distance [17, 26]; and the dimensionality d of the target space is typically low (i.e., $1 \leq d \leq 3$). The spatial model of preferences rests on solid theoretical foundations [7, 16, 19], and latent-space models could be viewed as implicitly extending these foundations beyond the political realm. However, while much *methodological* effort has focused on estimating voters' ideal points from real-world preference data [6, 10, 27, 29, 32], a growing body of work indicates that low-dimensional Euclidean space under the Euclidean norm provides a poor fit—both *conceptually* [18, 25, 33] and *empirically* [24, 34, 39, 43, 44]—even in political applications.

Motivation and Related Work. Given the importance of translating between ordinal and metric representations of preference information in practical applications, it is natural to ask under which conditions a given *preference profile* (i.e., a list of rankings over alternatives) permits what we call a *rank-preserving embedding*: a mapping of voters and alternatives into a normed real vector space that encodes the available ordinal preference information in the relationships between the pairwise distances. Toward answering this question, researchers have identified *necessary conditions* in terms of forbidden substructures [35, 41], and they have analyzed the computational complexity of recognizing when a preference profile admits a rank-preserving embedding into Euclidean space of a specified dimension [20, 28, 35]. They have also formulated *sufficient conditions* regarding the dimensionality of the embedding space for the special cases of the Euclidean [4] and Manhattan [9] norm. However, to our knowledge, nothing is known about the rank embeddability of preference profiles under *arbitrary norms*.

The dearth of rank-embeddability results under norms that are neither Euclidean nor Manhattan stands in stark contrast to the growing interest in concepts requiring the joint consideration of *all metric spaces*, such as *metric distortion* [see 2, 3, 5, 14]. Moreover, non-Euclidean norms have proved to be relevant across a variety of contexts—from *voting* [9, 15, 25, 36, 38] and *facility location* [22, 30, 31] to *algorithmic data analysis* [1, 11–13] and *recommender systems* [37]. Overall, a better understanding of the restrictions imposed on rank embeddability by more general norms appears desirable. Our work makes significant progress toward this goal.

Formal Setting and Intuition. We operate in a setting with m alternatives $A = \{a_1, \ldots, a_m\}$ (also known as candidates) and n voters $V = \{1, \ldots, n\}$. Given a set of alternatives A and a voter i, a *preference* of i is a weak order $\succsim_i$ on A, where $x \succsim_i y$ indicates that i weakly prefers alternative x to alternative y. A preference is *strict* if $x \sim_i y$ implies $x = y$, and it is *complete* if we have $x \succsim_i y$ or $y \succsim_i x$ for all $x \neq y$. We assume that preferences are strict and complete.

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/JWXO7778

Given alternatives A and voters V with preferences over A, the list $(\succsim_i)_{i=1}^n$ is said to be a *preference profile* $\mathcal{P}_{A,V}$ (or $\mathcal{P}$ for short). Seeking to characterize when a preference profile can be faithfully represented in a given metric space, for *real* metric spaces with *arbitrary* norms, we introduce a flexible notion of embeddability.

Definition 1 (Rank-preserving embedding, rank embeddability). *Given a preference profile $\mathcal{P}$, a dimension d, and a normed real vector space $(\mathbb{R}^d, \|\cdot\|)$, an assignment of coordinates $\mathbf{a}_j \in \mathbb{R}^d$ to alternatives $a_j \in A$ and coordinates $\mathbf{v}_i \in \mathbb{R}^d$ to voters $i \in V$ constitutes a* rank-preserving embedding *of $\mathcal{P}$ into $(\mathbb{R}^d, \|\cdot\|)$ if*

$$a_j \succsim_i a_k \iff \|\mathbf{v}_i - \mathbf{a}_j\| \leq \|\mathbf{v}_i - \mathbf{a}_k\| .$$

If there exists a rank-preserving embedding of $\mathcal{P}$ into $(\mathbb{R}^d, \|\cdot\|)$, $\mathcal{P}$ is said to rank-embed *into $(\mathbb{R}^d, \|\cdot\|)$.*

This definition generalizes the notions of *d-Euclidean* [4] and *d-Manhattan* [9] embeddings to arbitrary norms, and it can be further generalized to arbitrary metric spaces. Intuitively, Definition 1 requires that a coordinate assignment translates higher ranks in preference orderings to smaller distances in the embedding space. Since any two voters with identical preferences can be assigned identical coordinates, this implies that the rank embeddability of $\mathcal{P}$ into a given metric space only depends on the *distinct preferences* (i.e., *equivalence classes*) of voters. Therefore, we can assume w.l.o.g. that the preference profiles under study are *irreducible*, that is, no two voters i and j have identical preferences. Consequently, n can be interpreted as the number of distinct *voter types*.

2 MAIN CONTRIBUTIONS AND TECHNIQUES

Bogomolnaia and Laslier [4] proved that any preference profile with n voters and m alternatives rank-embeds into $(\mathbb{R}^d, \|\cdot\|_2)$ if $d \geq \min\{n, m-1\}$. Chen et al. [9] extended this result to $(\mathbb{R}^d, \|\cdot\|_1)$ using a very different approach, and they also treat the infinity norm using yet another construction. Leveraging our norm-independent notion of *rank embeddability* (Definition 1), we generalize these results to *arbitrary p-norms* for $p > 1$, streamlining the handling of the cases $p = \infty$ and $\infty > p > 1$ as a side effect:

Theorem 1 (Rank embeddability under p-norms). *Given m alternatives A and n voters V with preferences over these alternatives, a preference profile $\mathcal{P}_{A,V}$ rank-embeds into $(\mathbb{R}^d, \|\cdot\|_p)$, for all $1 \leq p \leq \infty$, if $d \geq \min\{n, m-1\}$.*

Separating the cases $d \geq n$ and $d \geq m-1$, our proofs are constructive and geometric, isolating only those features of the embedding classes introduced by Bogomolnaia and Laslier [4] that are strictly necessary to ensure rank preservation. Analyzing the asymptotic behavior of our main construction as $p \searrow 1$, we further elucidate why the 1-norm requires a fundamentally different construction.

We also explore how our results extend to arbitrary norms in low dimensions. In particular, we demonstrate that any preference profile with two voters rank-embeds into $(\mathbb{R}^2, \|\cdot\|)$ for *any* norm $\|\cdot\|$:

Theorem 2 (Rank embeddability for two [types of] voters under arbitrary norms). *Given m alternatives, let $\mathcal{P}$ be a preference profile featuring two (types of) voters. Then $\mathcal{P}$ rank-embeds into $(\mathbb{R}^2, \|\cdot\|)$ for any norm $\|\cdot\|$ on $\mathbb{R}^2$.*

Our proof relies on a geometric construction that, while not fully explicit, may still be of independent interest in the context of facility-location problems [e.g., 8, 30]. Unlike Bogomolnaia and Laslier [4], we do not need to resort to arithmetic calculations, which allows us to provide a much cleaner framework for reasoning about rank embeddability. All details can be found in the full version [45].

3 LIMITATIONS AND FUTURE WORK

Our contribution focuses on the *theoretical analysis* of spatial-preference models under arbitrary norms. Future work could systematically explore complementary approaches in this setting, such as *simulations* and *empirical studies*. On the theoretical side, while we made considerable progress toward a *full characterization* of rank embeddability under general norms, we stopped short of solving the general case. Hence, we explicitly pose the following conjecture.

Conjecture 1 (Rank-Embeddability Conjecture). *For $d \geq \min\{n, m-1\}$, any preference profile $\mathcal{P}$ with m alternatives and n voters can be rank-embedded into $(\mathbb{R}^d, \|\cdot\|)$, where $\|\cdot\|$ denotes any norm.*

Extending Theorem 1 to *polynomial norms* (including weighted p-norms) could be a natural next step toward full generalization. As the techniques used by Chen et al. [9] for $p = 1$ do not generalize to p-norms for $p > 1$, defining a *unified construction* that works for all p-norms also remains an open problem as well.

Beyond generalizing the upper bound on the dimensionality required to ensure rank embeddability, the *tightness of the bound* merits further investigation. Bogomolnaia and Laslier [4] show that in the presence of *indifferences* between alternatives (i.e., $a_j \sim_i a_l$ for some voter i and alternatives $j \neq l$), for any d, there always exist preference profiles with $d+1$ voters or $d+2$ alternatives that cannot be embedded into $\mathbb{R}^d$ under the Euclidean norm. However, their counterexample relies on a setting that requires at least one voter to be equidistant to multiple alternatives in the embedding space. This criterion makes the tightness of the bound significantly easier to address and does not cover, e.g., preference profiles that contain only strict inequalities or incomplete preferences. Hence, it would be valuable to understand if the bound is tight also in these scenarios.

Scrutinizing *low-dimensional profiles* under arbitrary norms could constitute another interesting avenue for future work. Escoffier et al. [21] provide some good combinatorial results for the amount of preference profiles that are Euclidean, Manhattan, or ℓ^∞ in low dimensions, while Chen et al. [9] highlight just how distinct the resulting embeddings can be. Against this background, studying which preference profiles rank-embed under *multiple norms* in a specific dimension might yield interesting insights. Given the generality of our result for $\mathbb{R}^2$ (Theorem 2), some of our techniques might prove useful for tackling this question.

Several other approaches also hold promise for approaching the general case (Conjecture 1). For instance, one could consider *approximations of balls* via polynomial norms. There are really robust uniformity results for any convex symmetric ball via polynomials—see, e.g., the works by Goodfellow et al. [23] or Totik [42]—, and we might be able to leverage those results for the higher-dimensional setting. While the corresponding constructions would not be explicit, any result establishing rank embeddability under arbitrary norms would provide a strong theoretical foundation for future work on problems related to spatial preferences.

REFERENCES

[1] Fateme Abbasi, Sandip Banerjee, Jaroslaw Byrka, Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, Dániel Marx, Roohani Sharma, and Joachim Spoerhase. 2023. Parameterized Approximation Schemes for Clustering with General Norm Objectives. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*. IEEE, 1377–1399. https://doi.org/10.1109/FOCS57990.2023.00085

[2] Ioannis Anagnostides, Dimitris Fotakis, and Panagiotis Patsilinakos. 2022. Metric-Distortion Bounds under Limited Information. *J. Artif. Intell. Res.* 74 (2022), 1449–1483. https://doi.org/10.1613/JAIR.1.13338

[3] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. 2018. Approximating Optimal Social Choice under Metric Preferences. *Artif. Intell.* 264 (2018), 27–51. https://doi.org/10.1016/J.ARTINT.2018.07.006

[4] Anna Bogomolnaia and Jean-François Laslier. 2007. Euclidean Preferences. *Journal of Mathematical Economics* 43, 2 (2007), 87–98. https://doi.org/10.1016/j.jmateco.2006.09.004

[5] Ioannis Caragiannis, Nisarg Shah, and Alexandros A. Voudouris. 2022. The Metric Distortion of Multiwinner Voting. *Artif. Intell.* 313 (2022), 103802. https://doi.org/10.1016/J.ARTINT.2022.103802

[6] Royce Carroll, Jeffrey B Lewis, James Lo, Keith T Poole, and Howard Rosenthal. 2013. The Structure of Utility in Spatial Models of Voting. *American Journal of Political Science* 57, 4 (2013), 1008–1028. https://doi.org/10.1111/ajps.12029

[7] Christopher P Chambers and Federico Echenique. 2020. Spherical Preferences. *Journal of Economic Theory* 189 (2020), 105086. https://doi.org/10.1016/j.jet.2020.105086

[8] Hau Chan, Aris Filos-Ratsikas, Bo Li, Minming Li, and Chenhao Wang. 2021. Mechanism Design for Facility Location Problems: A Survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*. ijcai.org, 4356–4365. https://doi.org/10.24963/IJCAI.2021/596

[9] Jiehua Chen, Martin Nöllenburg, Sofia Simola, Anaïs Villedieu, and Markus Wallinger. 2022. Multidimensional Manhattan Preferences. In *LATIN 2022: Theoretical Informatics - 15th Latin American Symposium, Guanajuato, Mexico, November 7-11, 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13568)*. Springer, 273–289. https://doi.org/10.1007/978-3-031-20624-5_17

[10] Joshua Clinton, Simon Jackman, and Douglas Rivers. 2004. The Statistical Analysis of Roll Call Data. *American Political Science Review* 98, 2 (2004), 355–370. https://doi.org/10.1017/S0003055404001194

[11] Vincent Cohen-Addad, Fabrizio Grandoni, Euiwoong Lee, and Chris Schwiegelshohn. 2023. Breaching the 2 LMP Approximation Barrier for Facility Location with Applications to k-Median. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*. SIAM, 940–986. https://doi.org/10.1137/1.9781611977554.CH37

[12] Vincent Cohen-Addad and Karthik C. S. 2019. Inapproximability of Clustering in Lp Metrics. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*. IEEE Computer Society, 519–539. https://doi.org/10.1109/FOCS.2019.00040

[13] Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. 2022. Johnson Coverage Hypothesis: Inapproximability of k-Means and k-Median in p-Metrics. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9-12, 2022*. SIAM, 1493–1530. https://doi.org/10.1137/1.9781611977073.63

[14] Soroush Ebadian, Daniel Halpern, and Evi Micha. 2024. Metric Distortion with Elicited Pairwise Comparisons. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*. ijcai.org, 2791–2798. https://doi.org/10.24963/ijcai.2024/309

[15] Daniel Eckert and Christian Klamler. 2010. An Equity-Efficiency Trade-Off in a Geometric Approach to Committee Selection. *European Journal of Political Economy* 26, 3 (2010), 386–391. https://doi.org/10.1016/j.ejpoleco.2009.11.009

[16] Jon X Eguia. 2011. Foundations of Spatial Preferences. *Journal of Mathematical Economics* 47, 2 (2011), 200–205. https://doi.org/10.1016/j.jmateco.2010.12.014

[17] Jon X Eguia. 2013. On the Spatial Representation of Preference Profiles. *Economic Theory* 52, 1 (2013), 103–128. https://doi.org/10.1007/s00199-011-0669-8

[18] Jon X Eguia et al. 2009. Utility Representations of Risk Neutral Preferences in Multiple Dimensions. *Quarterly Journal of Political Science* 4, 4 (2009), 379–385. https://doi.org/10.1561/100.00009059

[19] James M Enelow and Melvin J Hinich. 1984. *The Spatial Theory of Voting: An Introduction*. Cambridge University Press.

[20] Bruno Escoffier, Olivier Spanjaard, and Magdaléna Tydrichová. 2023. Algorithmic Recognition of 2-Euclidean Preferences. In *ECAI 2023*. IOS Press, 637–644. https://doi.org/10.3233/FAIA230326

[21] Bruno Escoffier, Olivier Spanjaard, and Magdaléna Tydrichová. 2024. Euclidean Preferences in the Plane under ℓ_1, ℓ_2 and ℓ_∞ Norms. *Soc. Choice Welf.* 63, 1 (2024), 125–169. https://doi.org/10.1007/S00355-024-01525-2

[22] Itai Feigenbaum, Jay Sethuraman, and Chun Ye. 2017. Approximately Optimal Mechanisms for Strategyproof Facility Location: Minimizing lp Norm of Costs. *Mathematics of Operations Research* 42, 2 (2017), 434–447. https://doi.org/10.1287/moor.2016.0810

[23] Ian J. Goodfellow, Patrick D. McDaniel, and Nicolas Papernot. 2018. Making Machine Learning Robust against Adversarial Inputs. *Commun. ACM* 61, 7 (2018), 56–66. https://doi.org/10.1145/3134599

[24] Marc Henry and Ismael Mourifié. 2013. Euclidean Revealed Preferences: Testing the Spatial Voting Model. *Journal of Applied Econometrics* 28, 4 (2013), 650–666. https://doi.org/10.1002/jae.1276

[25] Macartan Humphreys and Michael Laver. 2010. Spatial Models, Cognitive Metrics, and Majority Rule Equilibria. *British Journal of Political Science* 40, 1 (2010), 11–30. https://doi.org/10.1017/S0007123409990263

[26] Tasos Kalandrakis. 2010. Rationalizable Voting. *Theoretical Economics* 5, 1 (2010), 93–125. https://doi.org/10.3982/TE425

[27] In Song Kim, John Londregan, and Marc Ratkovic. 2018. Estimating Spatial Preferences from Votes and Text. *Political Analysis* 26, 2 (2018), 210–229. https://doi.org/10.1017/pan.2018.7

[28] Vicki Knoblauch. 2010. Recognizing One-Dimensional Euclidean Preference Profiles. *Journal of Mathematical Economics* 46, 1 (2010), 1–5. https://doi.org/10.1016/j.jmateco.2009.05.007

[29] Krishna K Ladha. 1991. A Spatial Model of Legislative Voting with Perceptual Error. *Public Choice* 68, 1 (1991), 151–174. https://doi.org/10.1007/BF00173825

[30] Richard C. Larson and Ghazala Sadiq. 1983. Facility Locations with the Manhattan Metric in the Presence of Barriers to Travel. *Oper. Res.* 31, 4 (1983), 652–669. https://doi.org/10.1287/OPRE.31.4.652

[31] Euiwoong Lee and Kijun Shin. 2025. Facility Location on High-Dimensional Euclidean Spaces. In *16th Innovations in Theoretical Computer Science Conference, ITCS 2025, January 7-10, 2025, New York, NY, USA (LIPIcs, Vol. 325)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 70:1–70:21. https://doi.org/10.4230/LIPICS.ITCS.2025.70

[32] Carolina Luque and Juan Sosa. 2025. Operationalizing Legislative Bodies: A Methodological and Empirical Perspective with a Bayesian Approach. *J. Amer. Statist. Assoc.* 120, 549 (2025), 572–583. https://doi.org/10.1080/01621459.2024.2413928

[33] Jeffrey Milyo. 2000. Logical Deficiencies in Spatial Models: A Constructive Critique. *Public Choice* 105, 3 (2000), 273–289. https://doi.org/10.1023/A:1005203611524

[34] Sergi Pardos-Prado and Elias Dinas. 2010. Systemic Polarisation and Spatial Voting. *European Journal of Political Research* 49, 6 (2010), 759–786. https://doi.org/10.1111/j.1475-6765.2010.01918.x

[35] Dominik Peters. 2017. Recognising Multidimensional Euclidean Preferences. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*. AAAI Press, 642–648. https://doi.org/10.1609/AAAI.V31I1.10616

[36] Hans Peters, Hans van der Stel, and Ton Storcken. 1993. Generalized Median Solutions, Strategy-Proofness and Strictly Convex Norms. *Zeitschrift für Operations Research* 38, 1 (1993), 19–53. https://doi.org/10.1007/BF01416005

[37] Cynthia Rudin. 2009. The P-Norm Push: A Simple Convex Ranking Algorithm that Concentrates at the Top of the List. *J. Mach. Learn. Res.* 10 (2009), 2233–2271. https://doi.org/10.5555/1577069.1755861

[38] Sooahn Shin, Johan Lim, and Jong Hee Park. 2025. ℓ^1-based Bayesian Ideal Point Model for Multidimensional Politics. *J. Amer. Statist. Assoc.* 120, 550 (2025), 631–644. https://doi.org/10.1080/01621459.2024.2425461

[39] Lukas F Stoetzer and Steffen Zittlau. 2015. Multidimensional Spatial Voting with Non-Separable Preferences. *Political Analysis* 23, 3 (2015), 415–428. https://doi.org/10.1093/pan/mpv013

[40] Donald E Stokes. 1963. Spatial Models of Party Competition. *American Political Science Review* 57, 2 (1963), 368–377. https://doi.org/10.2307/1952828

[41] Luke Thorburn, Maria Polukarov, and Carmine Ventre. 2023. Error in the Euclidean Preference Model. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence* (Macao, P.R.China) *(IJCAI '23)*. Article 322, 9 pages. https://doi.org/10.24963/ijcai.2023/322

[42] Vilmos Totik. 2013. Approximation by Homogeneous Polynomials. *J. Approx. Theory* 174 (2013), 192–205. https://doi.org/10.1016/J.JAT.2013.07.005

[43] Min Ye, Quan Li, and Kyle W Leiker. 2011. Evaluating Voter–Candidate Proximity in a Non-Euclidean Space. *Journal of Elections, Public Opinion & Parties* 21, 4 (2011), 497–521. https://doi.org/10.1080/17457289.2011.609619

[44] Xingchen Yu and Abel Rodriguez. 2021. Spatial Voting Models in Circular Spaces: A Case Study of the US House of Representatives. *The Annals of Applied Statistics* 15, 4 (2021), 1897–1922. https://doi.org/10.1214/21-AOAS1454

[45] Joshua Zeitlin and Corinna Coupette. 2025. Real Preferences Under Arbitrary Norms. arXiv:2508.09046 [econ.TH] https://arxiv.org/abs/2508.09046