THE SAGA OF KPR: THEORETICAL AND EXPERIMENTAL DEVELOPMENTS

KIRAN SHARMA,[1,*] ANAMIKA,[1,†] ANINDYA S. CHAKRABARTI,[2,‡] ANIRBAN CHAKRABORTI,[1,$] AND SUJOY CHAKRAVARTY[3,#]

In this article, we present a brief narration of the origin and the overview of the recent developments done on the Kolkata Paise Restaurant *(KPR) problem, which can serve as a prototype for a broader class of resource allocation problems in the presence of a large number of competing agents, typically studied using coordination and anti-coordination games. We discuss the KPR and its several extensions, as well as its applications in many economic and social phenomena. We end the article with some discussions on our ongoing experimental analysis of the same problem. We demonstrate that this provides an interesting picture of how people analyze* complex *situations, and design their strategies or react to them.*

Introduction

Chakrabarti et al. introduced the *Kolkata Paise Restaurant* problem (KPR henceforth), which was a multi-agent, multi-choice game[1,2]. The original motivation of the problem was to introduce the statistical analysis of a large-scale coordination game. The game theoretic solution concepts were introduced and it was shown that the agents (bots) could attain surprising level of coordination using some simple heuristics, even without any global coordinator. This approach was parallel to the one followed by Brian Arthur in his famous paper on El Farol bar problem (EFB henceforth)[3], which later gave birth to the *minority game* (MG henceforth) formulated by the physicists[4].

1   School of Computational and Integrative Sciences, Jawaharlal Nehru University, New Delhi-110067, India

2   Economics area, Indian Institute of Management, Ahmedabad, Gujarat, India

3   School of Social Sciences, Jawaharlal Nehru University, New Delhi-110067, India

*   e-mail : kiran34_sit@jnu.ac.in

†   e-mail : anamika@jnu.ac.in

‡   e-mail : anindyac@iima.ac.in

$ e-mail : anirban@jnu.ac.in

e-mail : sujoyc@gmail.com

The essential idea of Arthur was to show that heuristics based on induction, can be very useful for solving a game theoretic scenario. The KPR problem can be viewed as simply a straight-forward generalization of the EFB problem, yet mathematically tractable. However, as we would mention in this paper, notwithstanding the huge amount of research done on the MG problem[5], KPR provides a bigger laboratory to study many "complex" emergent phenomena. In recent times, it has been also used for operations research problems, which have enormous practical applications. For example, the taxi services (like Uber and Ola) solve a bilateral matching problem and KPR provides a starting point for analyzing such games[6]. For the sake of completeness, we state only two versions of the KPR problem here, though there are good resources for a much more complete description and discussion. Interested readers can refer to the following papers.[7,8,9]

We can define two versions of the KPR, one without ranking and the second with ranking. The first one is simpler and closer to a standard (anti-) coordination game, and so we begin with the one without ranking.

Game 1 Consider a set of *n* players and a set of *m* identical restaurants (*n = m* is the most interesting case). Each player has access to its own set of

information, which can be null (in principle), and has to secure access to one restaurant (win). Each of the players has to decide which restaurant to go, and if more than one player arrives at any restaurant, then only one of them is chosen *randomly* with equal probability.

Game 2 Identical to the above, except that the restaurants are now *ranked*, which is commonly agreed upon and known to all the players.

For the first game (and with the symmetric case, $n = m$), the Nash equilibrium could be simply achieved by assigning each of the agents to one restaurant only, such that all m restaurants are allocated and there are n winners. In the second case, it is certainly more complicated – imagine all the n agents being assigned to the m restaurants one-to-one, and solving the KPR problem. Since the restaurants are now ranked, there would be incentives for the person being assigned to lower rank restaurants to then deviate and choose upper ranked restaurants. This type of cases have been considered in details[9], and Chakrabarti et al.[1] provides conditions under which there exist pure strategy Nash equilibria[1]. Below, we briefly discuss the different facets of the KPR.

KPR at the Crossroads of Disciplines

There exists signicant amount of research done on coordination problems in multiple disciplines. Computer scientists have worked on parallel processing of jobs and tried to solve allocation problems by generating efficient algorithms. Statistical physics literature has dealt with traffic models, which again relate to individual choices leading to congestion and hence inefficiency. Obviously, from a game theoretic perspective, the apparently different scenario can be understood essentially as one *congestion* game. There are many other possible applications and interpretations of the game. Hence, we consider the KPR problem to represent a paradigm rather than a single, specific problem. In the following, we describe some directions the KPR problem has been studied and the links with existing literature across disciplines.

Heuristics to Solve a Complex Coordination Problem : A major part of the research on KPR has been to model it as a complex coordination problem and show two things: Firstly, simple and local heuristics can be useful for solving global optimization problem. Secondly, it can provide insights into how effective the different heuristics are to solve the problem. Imagine that this is a resource competition game and each successful match leads to

reproductive success. Then the strategy pool would start evolving favoring more successful strategies. Chakrabarti and Ghosh[10,11] considered a sequence of heuristics and also introduced limited form of learning and analyzed the relative successes. It turned out that simple rules of reinforcement learning with little use of any global or local information of the agents, could actually solve the allocation problem with very high efficiency.

Attaining Distributed Coordination : An important problem in designing efficient solvers is to find a way to distribute the load across different processors to attain parallel processing. One way to think about it, is to consider a central coordinator that would assign tasks across processors. However, in a multi-agent system with a very large number of agents, it may be prohibitively expensive to program such task (or resource) allocation *a priori* for all possible contingencies. Therefore, we may like to have *intelligent* agents who can strategize and discover the way to divide tasks by themselves.

Cigler and Faltings[12] considered this approach and suggested that a correlated equilibrium (proposed in the game theory literature by Aumann) could solve this problem; but they also recognized that such a device (which induces anti-coordination) basically acts as the central coordinator, and so this idea did not prove to be very useful. So they further proposed some learning rules, which are useful for solving this kind of problems. Hansen and Giauque[13] proposed dynamic programming algorithm as an alternate for task allocation problems. There are many other papers, which addressed similar questions[14,15,16].

Connections to Game Theory : The KPR is essentially a game theoretic problem, and so it is least surprising that economists would be interested in a similar paradigm. The original formulation of the KPR problem assumed that not only the agents have a ranking over the restaurants, but also the ranking is identically agreed upon and known to all agents. More explicitly, each and every agent agrees on which one is the better restaurant in any given pair of restaurants. A more complex version of the same problem was already known as the *Stable Marriage Problem*, where one assumed that there were two groups of members (identical in number), say men and women, who would marry each other[a]. Each man had a ranking of the women, and vice versa. The problem was to find the matches across members of the two groups such that there would be no two members who would rather break up and marry each other.

Gale and Shapley in 1962 provided an algorithm to solve this problem[17]. This problem has related variants like assignment problem, where a number of agents have to be

---

[a]   Same-sex marriage is not considered in the present context.

matched with the same number of tasks such that the matching is one-to-one. The catch is that every match has a cost of implementation and so the problem is to minimize the total cost of all matches. Such problems are of immense importance from a practical point of view. Roth and Peranson[18] provided a detailed algorithm along with applications for solving such allocation problems.

It is clear that although KPR deals with an allocation problem, there exist significant differences with respect to the above-mentioned approaches. One of the crucial differences is that in the above problems, it is assumed that there exists a central coordinator who allocates the task, whereas in KPR there exists none. So, we now move from the discussion on a *central planner* to a *decentralized* outcome.

From the game theoretic perspective, an important class of games known as *congestion games*, are close to the KPR problem. This class of games also assumes that players can choose any one among multiple options (roads) and if the same option is chosen by more than one player (leading to congestion) then that is bad for all of them. This can be mapped to the KPR problem (without ranking) upon simplification.

The closest problem was stated by Grenager *et al.*[19], as well as Alpern and Reyniers[20] who called it *dispersion game*. However, KPR with ranking has a more general structure. Below, we discuss the connection of the KPR with agent-based modelling literature and in particular, with the minority game literature.

Agent-based Models and Connections to MG

A variant of the EFB problem was proposed and analyzed by Challet and Zhang, which later on gained considerable popularity as the minority game (MG)[4,21-25]. Structure-wise, MG is a *N*-agents, 2-choice repeated game. Choices are defined over two restaurants are considered and the agents on the less-crowded side are considered to be the winners while the rest of them are declared losers. The primary aim of introducing the MG was to serve as simple a model of a financial market. Two choices are identified as 'buy' or 'sell' options executed by the trades. More buyers than sellers in any day would imply gain (positive payoff) for sellers and the opposite when there were more sellers than buyers. The total payoff of all agents could be treated as equivalent to the excess demand of the market, and be considered as the price return (directly proportional to the excess demand) of the market[21].

---

b   See Refs.[33] and[34] for meta-analyses of prisoners' dilemma and linear public goods games, respectively.

c   See Ref.[41] for a meta-analysis of trust game experiments.

Traders would learn from their own history of the game and alter their optimization strategy as the need be. The fluctuations of the players using different strategies in the steady state reflects the inherently dynamical and adaptive nature of the game. Agent-based modeling and Monte Carlo simulations, along with tools borrowed from statistical physics were used to show that the variability of fluctuation can be minimized in terms of memory size and number of players.

Physicists and economists became interested in related problems like the crowd versus anti-crowd movement of agents, broadly in the context of opinion formation. MG has served as prototype model for evolving and adaptive systems. Sysi-Aho *et al.*[26-29], in a sequence of papers, studied a modified version of the MG where the players were able to choose newer strategies using mechanisms of genetic crossover depending on their relative performances. This modification was useful to make the agent-based models applicable to many real-world systems especially with heterogeneous agents and non-linear interactions.

Could Laboratory Experiments Provide New Insight?

Numerous laboratory experiments have tested whether or not human agents play Nash equilibrium strategies in various different games, both one-shot and sequential. Some one-shot games with unique pure strategy Nash equilibria that have been tested in the laboratory include the prisoners' dilemma[30,31], voluntary contributions mechanism (VCM) games[33,34], auctions[35], Bertrand pricing game[36] and the travelers' dilemma[37]. A few sequential games, including the ultimatum bargaining game[38,39] and the trust game[40], have been studied in the laboratory. The results from these studies broadly establish that human players are not as *hyper-rational* as predicted by theory, and the laboratory outcome is often not the Nash equilibrium[4]. In almost all studies, *bounded rationality* as well as concerns for *equality*, affect human decision-making. Stahl and Wilson[42] introduced level-*k* rationality model that attempts to model different levels of rationality that a human player in a one-shot game may display. Models of other-regarding preferences that have been informed by results from laboratory experiments include those by Rabin[43], Fehr and Schmidt[44], and Bolton and Ockenfels[45].

The class of games that is particularly relevant for our study is coordination games. These are modeled in their normal form as one shot games with multiple Nash equilibria that can be Pareto rankable[46] or non-Pareto rankable[47,48]. In such games, an empirical question that interests game theorists and experimental economists is

whether agents manage to play the coordination outcome that realizes for them the highest level of joint profit, i.e., the welfare maximizing or efficient outcome. Given that a refinement is necessary to identify the more likely ones from this set of multiple equilibria, game theorists and experimentalists have developed possible selection criteria of equilibria and tested their empirical validity. Broadly speaking, the selection criteria that have been suggested either rely on structural properties of games or salience in the form of naturally occurring psychological labels. An important criterion relying on the structural properties of games involves the ideas of payoff and risk dominance was introduced by Harsanyi and Selten[49]. The theory of hierarchical decision-making that relies on the existence of focal points was first posited by Schelling[47] and tested by Mehta et al.[48]. Accordingly, naturally occurring labels that are of common knowledge can allow for greater coordination, particularly, in situations where there exist many non-Pareto rankable Nash equilibria.

Important experimental studies of coordination games with multiple Pareto ranked Nash equilibria includes the work of Van Huyck et al.[50], who study the minimum effort coordination game with seven Pareto rankable Nash equilibria with 14 to 16 subjects. They find on static repetition that subjects initially attempt to coordinate on the payoff dominant Nash equilibria but within a few periods only manage to converge to the risk dominant equilibrium. Cooper et al.[51] use a much simpler two player three strategy design, and in a similar fashion and that contrary to what Harsanyi and Selten[49] argue. Pareto dominance does not prove to be a robust predictor for equilibrium selection and dominated strategies are often not decision-irrelevant. Straub[52] reaches a conclusion similar to Cooper et al.[51] by using simple two by two coordination games.

Strong path dependence is seen in another work by Huyck et al.[53] in a median-action "continental divide" game, where there are two Pareto ranked equilibria. If in the initial rounds of a finitely repeated game session, participants are close to the high payoff equilibrium, then they tend to remain close to this equilibrium over all the rounds. A similar pattern is seen for sessions where choices remain close to the lower payoff equilibrium for the whole session. This history dependence in statically repeated one-shot games opens up the possibility for applying lagged behavioural models such as the reinforcement learning model[54] and the Experience Weighted Attraction (EWA) model[55].

An important goal of our KPR experiment[56] is to test the various strategies that are employed by incentivized human subjects in the complex situation where a solution cannot be deduced. There are a number of points that we would like to see: (i) Whether or not these strategies move the cohort of players closer to the efficient outcome; (ii) If there is any effect of experience and information feedback on coordination to the efficient outcome; (iii) The usage of learning models in order to model the decisions taken in the experiment, and (iv) The effect of external mechanisms that can potentially increase coordination efficiency and lead to better utilization.

In the experimental data collected so far, we do not see a signicant effect of information and experience in increasing the utilization ratio. In our experiment, we broadly observe a few behavioral types: (a) the noise traders – who are reactive and change their choices very frequently, (b) the completely stable agents – who stick to one strategy over all the periods, and (c) the intermediate type – who change their strategies a few times over the experiment, though not as much as the agents in case (a). We also check whether changes in strategy are influenced by the outcome in the previous period and find that going from failure to success in the previous period signicantly lowers the likelihood of a player changing his or her strategy in the current period. The results do not exactly match with similar simulation studies that were conducted before, and hence provide new insights. We do need to conduct more such experiments and understand how human beings *actually* strategize and adapt in "complex" situations.

Summary and Outlook

The KPR serves as a prototype model describing a general problem of competition over finite resources in a large population with a fully decentralized decision process, i.e., without any communication between the competitors. The repeated interactions and the associated learning mechanism lead to emergence of collective dynamics in the form of fluctuations in resource utilization. Additionally, in this competitive framework, the agents choose to be in the minority. Hence, they need to act as differently as possible through adaptation and modication of strategies, which would entail heterogeneous and non-equilibrium dynamics such as those found in many complex systems. As described above, the KPR encapsulates multiple facets of complexity theory and connects multiple disciplines. We hope that in the near future, this game will serve as a cornerstone for studying global optimization problem from localized active *smart* agents, in both theory and experiment.

References

1. A. S. Chakrabarti, B. K. Chakrabarti, A. Chatterjee, M. Mitra, The Kolkata Paise Restaurant problem and resource utilization, *Physica A*, **388**, 2420, (2009)

2. A. Ghosh, A. Chatterjee, M. Mitra, B. K. Chakrabarti, Statistics of the Kolkata paise restaurant problem, *New Journal of Physics,* **12**, 075033, (2010)

3. W. B. Arthur, Inductive reasoning and bounded rationality: The El Farol problem, *American Economic Review,* **84**, 40, (1994)

4. D. Challet, Y-C. Zhang, Emergence of cooperation and organization in an evolutionary game, *Physica A,* **246**, 407, (1997)

5. D. Challet, M. Marsili, Y-C. Zhang, Minority Games: Interacting Agents in Financial Markets, Oxford University Press, (2004)

6. P. Yang, K. Iyer, P. Frazier, Mean field equilibriua for resource competition in spatial settings, arxiv:1707.07919

7. A. Chakraborti, D. Challet, A. Chatterjee, M. Marsili, Y-C. Zhang, B. K. Chakrabarti, Statistical mechanics of competitive resource allocation using agent-based models, *Physics Reports,* **552,** 1, (2015)

8. S. Sinha, A. Chatterjee, A. Chakraborti, B. K. Chakrabarti, Econophysics: An Introduction, WileyVCH, Berlin, (2010)

9. B. K. Chakrabarti, A. Chatterjee, A. Ghosh, S. Mukherjee, B. Tamir, Econophysics of the Kolkata Restaurant Problem and Related Games: Classical and Quantum Strategies for Multi-agent, Multi-choice Repetitive Games, Springer, (2017)

10. D. Ghosh, A. S. Chakrabarti, Emergence of distributed coordination in the Kolkata paise restaurant problem with finite information, *Physica A,* **483**, 16, (2017)

11. A. S. Chakrabarti, D. Ghosh, Emergence of anti-coordination through reinforcement learning in generalized minority games, *Journal of Economic Interaction & Coordination*, https://doi.org/10.1007/s11403-017-0204-5, (2017)

12. L. Cigler, B. Faltings, Decentralized anti-coordination through multi-agents learning, *Journal of Articial Intelligence Research,* **47**, 441, (2013)

13. J. Hansen, W. Giauque, Task allocation in distributed processing systems, *Operations Research Letters,* **5**, 137, (1986)

14. W. Chu, L. Holloway, M. Lan, K. Efe, Task allocation in distributed data processing, *Computer,* **13**, 57, (1980)

15. K. Doty, P. McEntire, J. O'Reilly, Task allocation in a distributed computer system, *in Proceedings of IEEE INFOCOM*, **33**, (1982)

16. A. Galstyan, K. Czajkowski, K. Lerman, Resource allocation in the grid with learning agents, *Journal of Grid Computing,* **3**, 91, (2005)

17. D. Gale, L. S. Shapley, College admissions and the stability of marriage, *American Mathematical Monthly,* **69**, 9, (1962)

18. R. Alvin, E. Peranson, The redesign of the matching market for American physicians: Some engineering aspects of economic design, *American Economic Review,* **89(4)**, 756, (1999)

19. T. Grenager, R. Powers, Y. Shoham, Dispersion games: General definitions and some specific learning results, AAAI-02 Proceedings, (2002)

20. S. Alpern, D. J. Reyniers, Spatial dispersion as a dynamic coordination problem, *Theory and Decision*, **53**, 29, (2002)

21. D. Challet, M. Marsili, Y-C. Zhang, Minority Games, Oxford University Press, (2005)

22. D. Challet, Y-C. Zhang, *Physica A,* **256**, 514, (1998)

23. E. Moro, The Minority Game: an introductory Guide, in Advances in Condensed Matter and Statistical Physics, Eds., E. Korutcheva, R. Cuerno, *Nova Science Publishers*, (2004)

24. W. Kets, The minority game: An economics perspective, arXiv:0706.4432

25. A. Chakraborti, I. Muni Toke, M. Patriarca, F. Abergel, *Quantitative Finance,* **(11)7**, 991, (2011)

26. M. Sysi-Aho, A. Chakraborti, K. Kaski, Adaptation using hybridized genetic crossover strategies, *Physica A*, **322**, 701, (2003)

27. M. Sysi-Aho, A. Chakraborti, K. Kaski, Biology helps you to win a game, *Physica Scripta T,* **106**, 32, (2003)

28. M. Sysi-Aho, A. Chakraborti, K. Kaski, Intelligent minority game with genetic crossover strategies, *European Journal of Physics B,* **34**, 373, (2003)

29. M. Sysi-Aho, A. Chakraborti, K. Kaski, Searching for good strategies in adaptive minority games, *Physical Review E,* **69**, 036125, (2004)

30. M. Flood, Some Experimental Games, *Management Science,* **5**, 5, (1958)

31. R. Dawes, Social dilemmas, *Annual Review of Psychology,* **31**, 169, (1980)

32. R. Isaac, J. Walker, Group Size Hypotheses of Public Goods Provision: An Experimental Examination, *Quarterly Journal of Economics,* **103**, 179, (1988)

33. D. Sally, Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958-1992, *Rationality and Society,* **7**, 58, (1995)

34. J. Zelmer, Linear Public Goods Games: A Meta Analysis, *Experimental Economics,* **6(3)**, 299, (2003)

35. J. H. Kagel, D. Levin, Auctions: Experiments, in The New Palgrave Dictionary of Economics, Macmillan, (2008)

36. M. Dufwenberg, U. Gneezy, Price competition and market concentration: an experimental study, *International Journal of Industrial Organization*, **18**, 7, (2000)

37. C. Capra, J. K. Goeree, R. Gomez, C. A. Holt, Anomalous Behavior in a Traveler's Dilemma?, *American Economic Review,* **89(3)**, 678, (1999)

38. W. Guth, R. Schmittberger, B. Schwarze, An Experimental Analysis of Ultimatum Bargaining, *Journal of Economic Behaviour and Organization*, **3**, 367, (1982)

39. R. Forsythe, J. Horowitz, N. Savin, M. Sefton, Fairness in Simple Bargaining Experiments, *Games and Economic Behavior*, **6**, 347, (1994)

40. J. Berg, J. Dickhaut, K. McCabe, Trust, reciprocity, and social-history, *Games and Economic Behavior,* **10(1)**, 122, (1995)

41. N. Johnson, A. Mislin, Trust games: A meta-analysis, *Journal of Economic Psychology,* **32**, 865, (2011)

42. D. Stahl, P. Wilson, On Players Models of Other Players: Theory and Experimental Evidence, *Games and Economic Behaviour,* **10**, 218, (1995)

43. M. Rabin, Incorporating Fairness into Game Theory and Economics, *American Economic Review*, **83(5)**, 1281-1302, (1993)

44. E. Fehr, K. Schmidt, A Theory of Fairness, Competition, and Cooperation, *Quarterly Journal of Economics*, **114(3)**, 817, (1999)

45. G. Bolton, A. Ockenfels, ERC: A Theory of Equity, Reciprocity, and Competition, *American Economic Review*, **90(1)**, 163, (2000)

46. S. Bowles, Microeconomics: Behaviour, Institutions and Evolution, Oxford University Press, (2004)

47. T. Schelling, The Strategy of Conict, Harvard University Press, (1960)

48. J. Mehta, C. Starmer, R. Sugden, The Nature of Salience: An Experimental Investigation of Pure Coordination Games, *American Economic Review,* **84(3)**, 658, (1994)

49. J. Harsanyi, R. Selten, A General Theory of Equilibrium Selection in Games, MIT Press, (1988)

50. J. Van Huyck, R. Battalio, R. Beil, Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure, *American Economic Review*, **80(1)**, 234, (1990)

51. R. Cooper, D. V. Dejong, R. Forsythe, T. W. Ross, Selection Criteria in Coordination Games, *American Economic Review,* **80(1)**, 218, (1992)

52. P. G. Straub, Risk Dominance and Coordination Failures in Coordination Games, *Quarterly Review of Economics and Finance*, **35(4)**, 339, (2005)

53. J. Van Huyck, R. Battalio, T. Cook, Adaptive Behaviour and Coordination Failure, *Journal of Economic Behaviour and Organization*, **32(4)**, 483, (1997)

54. I. Erev, A. Roth, Predicting how people play games: Reinforcement learning in games with unique mixed strategy equilibria, *American Economic Review,* **88**, 848, (1998)

55. C. Camerer, T. Hua Ho, Experience Weighted Attraction Learning in normal Form Games, *Econometrica*, **67(4)**, 827, (1999)

56. K. Sharma, Anamika, A. S. Chakrabarti, A. Chakraborti, S. Chakravarty, *manuscript in preparation,* (2017)