Journal of Computer and System Sciences 93 (2018) 11–29

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

On discrete preferences and coordination ☆,☆☆

Flavio Chierichetti [a], Jon Kleinberg [b], Sigal Oren [c,*]

[a] Sapienza University of Rome, Italy
[b] Cornell University, United States
[c] Ben-Gurion University of the Negev, Israel

ARTICLE INFO

Article history:
Received 20 March 2017
Received in revised form 14 September 2017
Accepted 1 November 2017
Available online 22 November 2017

Keywords:
Algorithmic game theory
Price of stability

ABSTRACT

An active line of research has considered games played on networks in which payoffs depend on both a player's individual decision and the decisions of her neighbors. A basic question that has remained largely open is to consider games where the players' strategies come from a fixed, discrete set, and where players may have different preferences among the possible strategies. We develop a set of techniques for analyzing this class of games, which we refer to as *discrete preference games*. We parametrize the games by the relative extent to which a player takes into account the effect of her preferred strategy and the effect of her neighbors' strategies, allowing us to interpolate between network coordination games and unilateral decision-making. We focus on the efficiency of the best Nash equilibrium and provide conditions on when the optimal solution is also a Nash equilibrium.

1. Introduction

People often make decisions in settings where the outcome depends not only on their own choices, but also on the choices of the people they interact with. A natural model for such situations is to consider a game played on a graph that represents an underlying social network, where the nodes are the players. Each node's personal decision corresponds to selecting a strategy, and the node's payoff depends on the strategies chosen by itself and its neighbors in the graph [2,4,14].

Coordination and internal preferences A fundamental class of such games involves payoffs based on the interplay between *coordination* — each player has an incentive to match the strategies of his or her neighbors — and *internal preferences* — each player also has an intrinsic preference for certain strategies over others, independent of the desire to match what others are doing. Trade-offs of this type come up in a very broad collection of situations, and it is worth mentioning several that motivate our work here.

- In the context of opinion formation, a group of people or organizations might each possess different internal views, but they are willing to express or endorse a "compromise" opinion so as to be in closer alignment with their network neighbors.

☆ This work has been supported in part by a Simons Investigator Award, a Google Research Grant, an ARO MURI grant, and NSF grants IIS-0910664, CCF-0910940, IIS-1016099 and a Microsoft Research Fellowship.

☆☆ A preliminary version of this paper was published at EC'13.

* Corresponding author.
E-mail addresses: flavio@di.uniroma1.it (F. Chierichetti), kleinber@cs.cornell.edu (J. Kleinberg), orensi@cs.bgu.ac.il (S. Oren).

https://doi.org/10.1016/j.jcss.2017.11.002
0022-0000/

- Questions involving technological compatibility among firms tend to have this trade-off as a fundamental component: firms seek to coordinate on shared standards despite having internal cost structures that favor different solutions.
- Related to the previous example, a similar issue comes up in cooperative facility location problems, where firms have preferences for where to locate, but each firm also wants to locate near the firms with which it interacts.

A line of work beginning in the mathematical social sciences has considered versions of this question — often motivated by the first class of examples above, concerned with opinion formation — where the possible strategies correspond to a continuous space such as $\mathbb{R}^d$ [6,10]. This makes it possible for players to adopt arbitrarily fine-grained "average" strategies from among any set of options, and most of the dynamics and equilibrium properties of such models are driven by this type of averaging. In particular, dynamics based on repeated averaging have been shown in early work to exhibit nice convergence properties [6], and more recent work including by two of the authors has developed bounds on the relationship between equilibria and social optima [1].

Discrete preferences In many settings that exhibit a tension between coordination and individual preferences, however, there is no natural way to average among the available options. Instead, the alternatives are drawn from a fixed discrete set — for example, there is only a given set of available technologies for firms to choose among, or a fixed set of political candidates to endorse or vote for. On a much longer time scale, there is always the possibility that additional options could be created to interpolate between what's available, but on the time scale over which the strategic interaction takes place, the players must choose from among the discrete set of alternatives that is available.

Among a small fixed set of players, coordination with discrete preferences is at the heart of a long line of games in the economic theory literature — perhaps the most primitive example is the classic *Battle of the Sexes* game, based on a pedagogical story in which one member of a couple wants to see movie A while the other wants to see movie B, but both want to go to a movie together. This provides a very concrete illustration of a set of payoffs in which the (two) players have (i) conflicting internal preferences (A and B respectively), (ii) an incentive to arrive at a compromise, and (iii) no way to "average" between the available options.

But essentially nothing is known about the properties of the games that arise when we consider such a payoff structure in a network context. Even the direct generalization of Battle of the Sexes (BoS) to a graph is more or less unexplored in this sense — each node plays a copy of BoS on each of its incident edges, choosing a single strategy A or B for use in all copies, incurring a cost from miscoordination with neighbors and an additional fixed cost when the node's choice differs from its inherent preference. Indeed, as some evidence of the complexity of even this formulation, note that the version in which each node has an intrinsic preference for A is equivalent to the standard network coordination game, which already exhibits rich graph-theoretic structure [14]. And beyond this, of course, lies the prospect of such games with larger and more involved strategy sets.

Formalizing discrete preference games In this paper, we develop a set of techniques for analyzing this type of discrete preference games on a network, and establish tight bounds on the efficiency of the best Nash equilibrium[1] for several important families of such games.

To formulate a general model for this type of game, we start with an undirected graph $G = (V, E)$ representing the network on the players, and an underlying finite set L of strategies. Each player $i \in V$ has a *preferred strategy* $s_i \in L$, which is what i would choose in the absence of any other players. Finally, there is a metric $d(\cdot,\cdot)$ on the strategy set L — that is, a distance function satisfying (i) $d(i,i) = 0$ for all i, (ii) $d(i,j) = d(j,i)$ for all i, j, and (iii) $d(i,j) \le d(i,k) + d(k,j)$ for all i, j and k. For $i, j \in L$, the distance $d(i,j)$ intuitively measures how "different" i and j are as choices; players want to avoid choosing strategies that are at large distance from either their own internal preference or from the strategies chosen by their neighbors.

Each player's objective is to minimize her cost: for a fixed parameter $\alpha \in [0,1)$, the cost to player i when players choose the strategy vector $z = \langle z_j : j \in V \rangle$ is

$$c_i(z) = \alpha \cdot d(s_i, z_i) + \sum_{j \in N(i)} (1-\alpha) \cdot d(z_i, z_j),$$

where $N(i)$ is the set of neighbors of i in G. The parameter α essentially controls the extent to which players are more concerned with their preferred strategies or their network neighbors; we will see that the behavior of the game can undergo qualitative changes as we vary α.

We say that the above formulation defines a *discrete preference game*. As standard in game theory, we will be interested in Nash equilibria of this game. These are strategy profiles in which no player can reduce its cost by deviating to a different strategy. Note that the network version of Battle of the Sexes described earlier is essentially the special case in which $|L| = 2$, and network coordination games are the special case in which $|L| = 2$ and $\alpha = 0$, since then players are only concerned with matching their neighbors. The case in which $d(\cdot,\cdot)$ is the distance metric among nodes on a path is also interesting to

[1] Formally, we prove bounds on the price of stability which we later define.