Forming K-Player games from N players using Collective Belief Aggregation

Kshanti A. Greene
Social Logic Institute
Albuquerque, New Mexico
greene@sociallogic.org

Joe M. Kniss and George F. Luger
Department of Computer Science
University of New Mexico
jmk,luger@cs.unm.edu

Abstract

Graphical games reduce the complexity of finding Nash equilibria in N-Player games by having each individual act based on a small number of neighbors' strategies. However, there are many situations (such as in political decision-making) in which one is not interested in how one's neighbor acts, but in the portion of the population that takes on a particular action. We outline an approach that enables one to find Nash equilibria in large populations by partitioning a population into *collectives* such that all members of a collective will take on a particular action given the distribution over actions of the rest of the population [1]. The utilities or expected utilities of each member in a collective are then aggregated to form a *super agent* that represents the collective in a K-player game, where K is the number of collectives. K is derived from the number of action choices and will typically be much smaller than N in a large group. When K is small one can find Nash equilibria for the population efficiently even when the graphical game is fully connected.

1 Introduction

We introduce an approach to enable efficient discovery of Nash equilibria in an N-Player game in which N is large. Our approach extends a Bayesian belief aggregation approach we developed to represent diverging beliefs in a population [Greene et al., 2010, Greene, 2010]. Bayesian belief aggregation is the process of forming a consensus model from the probabilistic beliefs of multiple individuals. Preference aggregation attempts to find an optimal solution for a population considering each individual's beliefs, desires and objectives. Belief and preference aggregation approaches that form a single consensus average away any diversity in a population. In the process they may fail to uphold a set of mathematical properties for rational aggregation defined by social choice theorists. Our aggregation approach maintains the diversity of a population and allows the competitive aspects of a situation to emerge. Each individual is represented by their beliefs and/or utilities. The population is separated into *collectives* whose members agree on the relatively likelihood or desirability of the possible outcomes of a situation. An aggregate for each collective can then be computed such that the aggregate upholds the rationality properties. Game theoretic analysis can then applied using "super-agents" that represent each collective as the game players. In this manner, we can find the Nash equilibria of a population of N people by partitioning the population into K collectives and forming a game of K players.

The motivation for this research is to utilize mathematical models to improve democratic decision-making. For example, we see promise in leveraging the rapid spread of information in popular social-networking tools to elicit opinions that form computational models for collective decision-making. These collective intelligence models can communicate the diverse ideas, beliefs and preferences of individual stakeholders to decision-makers so that community and business representatives can form policy that best serves their constituency. In this manner, these models will enable individuals to have a direct influence in the social, economic, and political decisions that affect them. In addition, game theoretic analysis can enable people to visualize how their actions affect their own circumstances and their environment, taking into consideration the simultaneous actions and goals of other community members.

This paper is organized as follows. In Section 2 we

[1] The research described in this article is in progress. Formalization of the approach is in development.

summarize three areas of research that address combining beliefs and preferences to form consensus models and game theory in large populations. We then summarize our belief aggregation approach in Section 3. Section 4 demonstrates how K-Player games are formed from N players and introduces a property for aggregation in games that maintains consistent Nash equilibria. We show that upholding this property depends on how collectives are formed. We then discuss the derivation of K and mixed strategy games.

2 Background

The research discussed in this paper combines three areas of research; Bayesian belief aggregation, social choice theory and game theory.

2.1 Bayesian Belief Aggregation

Belief aggregation is the process of combining probability estimates to form a *consensus* model from multiple human or software agents. Matzkevich and Abramson [Matzkevich and Abramson, 1992] cited two different approaches to belief aggregation. The first was called *posterior compromise*, which combines the beliefs after the network and probabilities have been defined and a query has been made. In other words, one would query separate networks and then combine the result. The authors introduced their alternative approach called *prior compromise* that instead found a consensus network *before* inference was done to determine the result of a query. This approach would involve fusing together networks that may also have different structure, called topological fusion. Once networks were fused, they combined the beliefs on local relationships using an approach called *family aggregation* [Pennock and Wellman, 1999].

Bayesian belief aggregation uses *opinion pool* functions whose output is a numeric result of the combination of a number of inputs. An opinion pool function is a mathematical function to form a single aggregate value from multiple beliefs. Mathematically, $P_0 = f(P_1, P_2, \ldots, P_n)$ where each P_i is the probability estimation from the i^{th} contributor given N contributors. P_0 is the *consensus* estimation. The two most commonly used opinion pools are the linear opinion pool (LinOP) and the logarithmic opinion pool (LogOP) [Pennock and Wellman, 1999].

Pennock and Wellman introduced a market-based belief elicitation and aggregation approach, which is intended to move individuals towards consensus. This approach requires that individuals back up their beliefs by buying and selling *stocks* that indicate their confidence in an event occurring [Pennock and Wellman, 2005]. The consensus value is determined by the resulting stock price. While this approach may improve accuracy when all agents have the same risk tolerance, this case is highly unlikely. In general a market based approach increases the subjectivity of the result as each individual has unequal desire to make a bet on their beliefs. In addition, the assumption that individuals' beliefs will move towards consensus by placing a financial risk on them ignores other non-market based factors that contribute to people's beliefs. In particular, one's background and experience play a strong role in political, religious and other subjective beliefs.

2.2 Social Choice Theory

Social choice theory, also called social welfare theory, is a branch of research that has involved researchers in voting theory, economics and statistics. Social choice theory analyzes the manner in which one can determine a *social choice*, or collective decision based on the opinions and preferences of a group of individuals. The area of research was launched by economist Kenneth Arrow's when he introduced his rationality properties for combining preferences and theorems on the limitations of finding a social choice [Arrow, 1950]. Many researchers followed to analyze and expand upon his findings in deterministic and Bayesian environments [Hylland and Zeckhauser, 1979, Feldman and Serrano, 2007].

Social choice and Bayesian theorists have stated that it is not possible to combine, or *aggregate* arbitrary beliefs or preferences to build a consensus model that conforms to a set of mathematical principles of rationality for shared preference [Arrow, 1950, Feldman and Serrano, 2007, Seidenfeld et al., 1989]. In particular, the economist Kenneth Arrow developed a theorem that states that one cannot aggregate arbitrary votes when there are three or more options to choose from without the possibility of violating at least one of the properties in Figure 1 for *rational* aggregation [Arrow, 1950].

2.2.1 N-Player Games

[Kearns et al., 2001] demonstrates an approach to discover the Nash equilibrium solutions in an n-player game in which n may be large. Instead of each player sharing their strategy with all other $n-1$ players, they share their strategy with their k neighbors. However, there are many situations in which each player only needs to know the number of players that are going to take on a particular strategy. Our approach more compactly represents the strategies of individuals by grouping them according to their strategies and con-

Given a society of interest, S:

1. *Universal Domain (UDP)*: All preference orders are allowed
2. *Completeness (CP)*: Social choice function returns an order that includes all relevant alternatives
3. *Transitivity (TP)*: if S prefers A to B and B to C then S prefers A to C (also replace "prefers" to "is indifferent to")
4. *Pareto optimality (should be at least weekly Pareto optimal)*:
   (a) Weak Pareto principle (WP). For all x and y, if xP_iy for all i, then xPy:
   (b) Strong Pareto principle (SP). For all x and y, if xR_iy for all i, and xP_iy for some i, then xPy:
5. *Independence of irrelevant alternatives: (IIP)* A society's preference order over a subset of options should be the same as the order over the whole set of options.
6. *Non-dictatorship and non-imposition (NDIP)*. There is no dictator. Individual i is a dictator if, $\forall x$ and $\forall y$, $xP_iy \rightarrow xPy$. Non-imposition means that no order has been pre-determined for any individual.

Figure 1: Rationality properties defined by Kenneth Arrow [Arrow, 1950, Arrow, 1963, Feldman and Serrano, 2007]

sidering the size of each collective as part of its strategy. [Daskalakis et al., 2006] shows that graphical N player games with k neighbors is in the same complexity class as a k player game. In our continued research we aim to show that the same efficiency benefits can be achieved without restricting each individual to share strategies with a small group of "neighbors."

3 Collective Belief Models

The following definitions form the foundation for our aggregation approach.

Definition 1 *Rank Order: A rank order* R *is a partial order* $\rho_1 \leq \rho_2 \leq ... \leq \rho_r$ *over a set of options* O *containing* r *options* $\mathrm{o_1..o_r}$*, where* ρ_i *is some* o_j*. A rank order* R *over* O *is an order such an item* ρ_i *is preferred to (or indifferent to)* ρ_j *if and only if* ρ_i *is before* ρ_j *in the order.*

For example, given $O = \{A, B, C\}$, if B is preferred or indifferent to C, and C is preferred or indifferent to A, then the rank order over O is $R = B \leq C \leq A$, (can also written as BCA). A *strict* preference for ρ_i over ρ_j is one such that $\rho_i < \rho_j$.

The term "collective" is often used in the study of society and social behavior and has many different sociological definitions. However, since no rigorous mathematical definition exists, we define a collective as a group such that a specific generalization of the group holds for all members of the group. In set theory, this is simply a subset of a set, with the generalization being the property that defines the subset.

Definition 2 *Collective: A collective* C *w.r.t. a property* A *is a subset of a population* P $(\mathrm{C} \subseteq \mathrm{P})$ *s.t.* A *holds for all members of* C*. If* A *holds for an individual* $\mathrm{p} \subseteq \mathrm{P}$ *then* p *is a member of* C*. A null set* $\emptyset$ *with respect to property* A *indicates that the generalization does not hold for any member of* P*.*

Arrow's theorem shows that there is no generalization that can be made about an arbitrary population based only on their preferences [Arrow, 1950, Feldman and Serrano, 2007]. However, if a group of individuals happens to have an identical rank order of their preferences, then we can easily make a generalization about that group based on this rank order. If the options are taken from a discrete set of values, and an individual i's rank order R_i is identical to another individual j's rank order R_j then a generalization R_0 can be made about the individuals i and j, such that $R_0 = R_i = R_j$. Thus, if a group of individuals G shares a rank order R over a set of discrete valued options O then R can define a collective C.

Definition 3 *Rank Order Collective: If* $\mathrm{R_j}$ *is a rank order over a set of options* O*, and* $\mathrm{C_j}$ *is a subset of a population* P *s.t.* $\forall \mathrm{p} \in \mathrm{C_j}$, $\mathrm{R_p} = \mathrm{R_j}$ *and* $\forall \mathrm{q} \notin \mathrm{C_j}$, $\mathrm{R_q} \neq \mathrm{R_j}$*, then* $\mathrm{C_j}$ *is a collective defined by the property* $\mathrm{R_j}$ *and is called a* rank order collective.

We will now map Bayesian outcomes to the rank order concept. We will use the terms *variable* and *node* interchangeably when referring to Bayesian networks. Each variable X_i in a Bayesian network has a number of possible values $(\{x_{i1}, x_{i2}, ..., x_{ir}\})$, where r is the arity of X_i. The posterior probability of a variable in a Bayesian network being a specific value (x_{ij}) is derived through inference [Pearl, 1988]. Given the posterior probabilities of the values of a variable, there will be an order from *most likely* to *least likely* for these possible values. For example, given a binary variable X, with a probability distribution $P(X = T) = 0.25, P(X = F) = 0.75$, the order of values is FT. This order is analogous to the rank order in Definition 1. In the case of a

Bayesian decision network, the result of inference is a set of *expected* utilities for the possible decision options [Howard and Matheson, 1984, Shachter, 1986]. The decision options can be ranked by order of *highest expected utility* to *lowest expected utility*, or *best* to *worst* option. Given a Bayesian network, the rank order can be determined for an arbitrary variable or decision in the network. A rank order can also be determined from a probability distribution over a set of outcomes or the expected utilities of a set of options.

Definition 4 *Bayesian Rank Order: A Bayesian rank order* R^* *with respect to a random variable* X *is a rank order over the posterior probabilities* $(\mathrm{P}(\mathrm{X}=\mathrm{x}_1),\mathrm{P}(\mathrm{X}=\mathrm{x}_2),...,\mathrm{P}(\mathrm{X}=\mathrm{x}_\mathrm{r}))$ *of the values of* X*. Likewise, given a decision* D *with* r *decision options* $\{\mathrm{d}_1,\mathrm{d}_2,...,\mathrm{d}_\mathrm{r}\}$*,* R^* *is the rank order of the expected utilities* $(\mathrm{EU}(\mathrm{d}_1),\mathrm{EU}(\mathrm{d}_2),...,\mathrm{EU}(\mathrm{d}_\mathrm{r}))$ *of the options of* D*, where* $\mathrm{EU}(\mathrm{d}_\mathrm{i})$ *is the expected utility of decision option* d_i*.*

It is trivial to see that using symbolic preference orders for which people simply provide a rank order over a number of options, grouping people according to their rank order would result in a consensus of the same rank order. However, when each individual's rank order is determined from their quantitative beliefs or expected utilities, it is not as obvious that this is the case. In [Greene, 2010] we showed that an aggregate of $m > 1$ equivalent Bayesian rank orders R over the probabilities of a variable X (or the expected utilities of a decision D), results in the same rank order, R. This was illustrated in [Greene, 2010] by showing that the mean of $m > 1$ arbitrary sets of n ordered values in $\mathbb{R}$ results in an ordered set of values.

We have defined the *Bayesian rank order* as the relative likelihood of a probabilistic variable's possible values or the relative expected utility of a set of decision options. The *collective belief* of a collective is the aggregate of the collective members' *Bayesian rank orders.*

Definition 5 *Collective Belief: If* R^* *is a Bayesian rank order over the probability distribution of a variable* X*, (or the expected utilities of a decision* D*), and* C *is a rank order collective s.t.* $\forall \mathrm{c}_\mathrm{i} \in \mathrm{C}, |\mathrm{C}| = \mathrm{k},$*, the Bayesian rank order of* c_i *is* R^**, then the* collective belief ϕ *of* C *is the aggregate of the* k *probability distributions (or expected utilities) supplied by the members of* C*.*

Definition 6 *Rational Social Choice: A collective belief* ϕ *is a rational social choice (RSC) if it conforms to the properties for rational belief aggregation defined in Figure 1.*

Definition 7 *Partition: A partition* T *of a population* P *is a set of collectives such that all individuals in* P *are in exactly one collective* C_j*.*

Definition 8 *Collective Belief Model: A collective belief model* $\mathrm{CBM} = (\mathrm{T}, \Phi)$ *is composed of a partition* T *containing* k *collectives* $\mathrm{C}_\mathrm{j} \in \mathrm{T}$*, and each collective's collective belief* $\phi_\mathrm{j} \in \Phi$ *for a given set of options* O*.*

3.1 Collective Choice Function

The following social choice function forms a partition and its collectives for a population, based on probability distributions or expected utilities provided by each member of the population. The *collective* choice function is defined as follows.

Definition 9 *Collective Choice Function: Given a population* P *and a set of* r *options* O*, over which each individual* i *in* P *provides a rank order* R_i*:*

i *Separate a population* P *into* m *groups, each group representing a unique ordering* R_j*, such that each individual* i *in group* G_j *provides a rank order* $\mathrm{R}_\mathrm{i} = \mathrm{R}_\mathrm{j}$ *over the* r *options in* O*.*

ii *Each group* G_j *becomes a collective* C_j *defined by its collective rank order* R_j*.*

iii *In a Bayesian environment, use an opinion pool function (e.g. arithmetic or geometric mean) to compute each collective* C_j*'s collective belief.*

[Greene, 2010] demonstrated that the *collective choice function* upholds the *RSC* properties by creating a partition T, such that each property holds for each collective in T. In summary, the properties are upheld for a collective because all members of a collective share the same rank order over a set of options.

4 Nash Equilibrium in K-Player Games

This section will first demonstrate how to apply the collective choice function in a game of pure strategy to find Nash equilibria using a K-Player extension of the Battle of the Sexes game. We show that aggregating utilities within a collective results in the same Nash equilibria as the N-player game. We then show that a slight variation in how the collectives are formed will result in a different Nash equilibria than an N-player game within the same population. Finally, a sketch of how the approach would be applied in mixed strategy games using a Bayesian rank order that considers the probability of action is given.

	LWE	LWM	LWW	WLE	WLM	WLW
LWE	**2,1**	X	X	0,0	X	X
LWM	X	**2.6,1.5**	X	X	2.4,0.25	X
LWW	X	X	2,1	X	X	0,2
WLE	0,0	X	X	**1,2**	X	X
WLM	X	0.2,0	X	X	0.8,6.5	X
WLW	X	X	0,0	X	X	**1,2**

Table 1: Normal form of Battle of the Sexes game, the mens' collective is in rows and womens' collective is in columns

4.1 Normal Form

K-Player games are formed by partitioning a population of N people into K collectives, each of which becomes a player. In K-Player games the strategy of each player includes the player's action and their *weight*, which is the percentage of the population that is in that collective. The actual percentage will be discretized into bins. In the following example three bins are formed: $\{< 50\%, 50\%, > 50\%\}$, meaning that the collectives are either equal in size or one collective has a simple majority.

We extend the Battle of the Sexes game to consider a population of men and women who are trying to decide on a movie [Leyton-Brown and Shoham, 2008]. The movie options are "Lethal Weapon" (LW) and "Wondrous Love" (WL). In our version of the game there are three possibilities for population distribution: the number of men and women are equal, there are more men than women, and there are more women than men. Given their druthers, the men would prefer to see LW while the women would prefer WL. However, similar to the two-agent game in which the man and woman would prefer to go to the same movie than separate movies, all of the people would prefer to go to the same movie as the majority.

We now show the normal form of a full game between two collectives. Suppose for illustration that we have partitioned the population into a male collective and a female collective. We then aggregated each collective's utilities for movies given the other collective's movie selection and the relative size of the collectives. The normal form of the game is shown in Table 1, in which the cells labeled with LWE and WLE contain the groups' utilities given the groups are equal size, cells labeled with LWM and WLM contain the utilities given that the men have the majority and cells labeled with LWW and WLW contain the utilities given that the women have the majority. For example, the cell (LWM, LWM) (row,col) contains the mens' and womens' aggregate utilities for seeing "Lethal Weapon" given that the men have the majority. While this normal form shows the strategy that includes the relative size of the groups, the only actions that will occur in reality are those in which the relative size is the same. Therefore the cell (LWM, LWW) and all other non-applicable combinations are marked with an 'X'. All Nash equilibria are marked in bold. When the group sizes are equal, there are two equilibria. However if the men or women have a majority, then the only equilibrium favors the majority's preference.

Men	LW,LW	LW,WL	WL,LW	WL,WL
Adam	2	2	0	1
Bob	3	2	0	0
Charlie	2	3	0	1
David	2	2	0	1
Edgar	4	3	1	1
Women	LW,LW	LW,WL	WL,LW	WL,WL
Alice	1	0	0	2
Betty	2	0	0	2
Christine	1	0	0	2
Diane	2	1	0	20

Table 2: Utilities of each man and woman given the actions of the women/men. Column LW,WL refers to the men choosing LW and women choosing WL.

4.2 Forming a K-Player Game from N players

Since the relative size strategies are mutually exclusive it is not necessary to combine them in a normal form. One can instead create a normal form for each possible situation. The remainder of this section focuses on when the men have the majority. We now show how a K-Player game is formed. In this case $K = 2$. Section 4.5 discusses how K is derived.

Table 2 shows the utilities of each person given the actions of the others and the fact that the men have the majority. All men in Table 2 have the same rank order over the movie options given the strategy of the women, and all women have the same rank order (which is different from the mens'). The mens' rank order is represented by $[LW, WL]|LW, [LW, WL]|WL$, meaning that they strictly prefer LW when the women choose LW and they strictly prefer LW when the women choose WL. The women are represented by $[LW, WL]|LW, [WL, LW]|WL$, meaning that they prefer LW when the men choose LW and they prefer WL when the men choose WL.

A person's actions can be determined from their rank order. In other words, they will always act the same given their ranked order regardless of the actual value of their utilities.

Proposition: *Assuming a person always takes their preferred action and there is a strict preference over the possible actions, a person's actions can be determined by the rank order over the actions given the pure strategies of the remaining players.*

Proof: The rank order states which action is preferred given the action of the other players. Therefore, if the action of the other players is known (as it is with pure strategy) the person will select the action with the highest rank. By the definition of rank order (Def. 1) a strictly preferred action will always be the highest rank.

If a person's actions can be determined by their rank order, regardless of the value of their utilities, then by the same logic, a *different* person with the same rank order will act the same in the identical situation.

Extension of previous proof: If any two people, a and b have the same rank order, $R_a = R_b$, over a set of actions given the actions of the other players, then a and b will select the same action given the actions of the remaining players.

As Section 3 demonstrated, if multiple individuals have the same rank order, then they can form a collective and their utilities can be aggregated to find the collective belief (or in this case collective utilities). Therefore, since all the men have the same rank order in Table 2 they can form a "mens' " collective and the women can form a "womens' " collective. In fact, the utilities in Table 1 for the male majority strategy are the collective beliefs (found with arithmetic mean) of the groups in Table 2.

4.3 Consistent Nash Equilibria

We now show that the set of Nash equilibria of the K-Player game as derived in Section 4.2 is the same as an N-Player game with the players in Table 2. For this example we drop the assumption that the men have higher weight and consider each man and woman as an individual. An N-player graphical game is formed such that each player is represented by a node. An undirected edge between a node representing a man, m and a node representing a woman, w means that m will make his movie selection based on w's movie selection and vice versa. Given any pairing of the men and women from Table 2, the Nash equilibrium will always be LW, LW (both the man and woman see "Lethal Weapon"). Again, this is because all men have the same rank order over the movie options given the womens' strategy, and vice versa.

In the spirit of Arrow's properties for rational social choice [Arrow, 1950, Greene, 2010], we introduce a property for consistent K-player games given N individuals.

Figure 2: A graph representing an N-player graphical game containing the players from Table 2. All players can be placed in this game with any of the men in an M node and any woman in any W node.

Property 1: Nash Equilibria in K-player games *The set of Nash equilibria $\overrightarrow{P_K}$ in a K-player game must be equivalent to the set of Nash equilibria $\overrightarrow{P_N}$ in an N-player game formed by selecting one individual at random from each of the K collectives. Any combination of players selected at random from the K collectives should result in $\overrightarrow{P_K} = \overrightarrow{P_N}$.*

Repeatedly selecting an individual at random from each collective duplicates an N-player graphical game as described in [Kearns et al., 2001] in which a tree is formed with each individual represented by a node. In our version, each person from collective k_i is connected by a edge to one or more people from each other collective $k_j, i \neq j$. Figure 2 shows a partial tree for the Battle of the Sexes game. Nodes marked W are women and M are men. Any assignment of nodes to the men and women in Table 2 would result in the same Nash equilibrium for the players.

4.4 Inconsistent Nash Equilibria

If a group is formed from individuals that differ in their rank order over the movie options given the movie option of another other group, the aggregate can result in a new Nash equilibrium. For example, if, instead of forming a mens' collective and a womens' collective we form two groups as follows: $G = \{Adam, Bob, Charlie, Diane, Edgar\}$ and $F = \{Alice, Betty, Christine, David\}$, the normal

	LW	WL
LW	**2.6,1.5**	2.2, 0.5
WL	0.2,0	**4.6, 1.75**

Table 3: Normal form of game with groups G (rows) and F (columns). Notice that this grouping forms a Nash equilibrium (WL,WL) that is not seen with the mens' and womens' collective.

form and aggregate utilities are shown in Table 3. Given this grouping, a new Nash equilibrium (A sees WL, B sees WL) arises. This is inconsistent with Property 1. We conclude that maintaining rank order within a group *may* be required to guarantee Property 1 when aggregating groups to form a K player game, although further investigation is needed to confirm this insight.

4.5 Deriving K

The value of K is not dependent on the size of the population (N). Instead, it is based on two parameters: the number of actions to choose from (A) and the number of bins over $[0, 1]$ the distribution over actions is discretized into when considering the strategies of the other collectives (B). The previous examples focused on a simple majority, for which there could be three sets of actions for each person: what action to take if both groups are equal, what action to take if men have the majority and what action to take if women have the majority. If we consider the strategies for all situations, then $B = 3$. Since we focused primarily on one situation in which men had the majority, we let $B = 1$. The following equation shows the formula for deriving K.

$$K = A!B \tag{1}$$

Derivation of A is described in [Greene, 2010]. In summary, it is the number of permutations of the set of A actions, equivalent to the number of possible rank orderings over the A actions. When the percentage of the population taking on each action is considered in the strategy, then there can be any combination of rank orderings *given* the distribution over the actions of the population. In the previous example in which $B = 1$ the distribution was always "men have the majority".

Since A is typically small and one is often only concerned in a simple majority, then K will often remain very small. As A or B grows, then we will only see improvements in efficiency if N is very large. However, we reiterate that since K is not dependent on N, then improvements in efficiency can be seen even when N grows to include (for example) the entire voting population of a country.

4.6 Mixed Strategies

This paper has thus far only discussed pure strategies, in which the strategies of each player is known. In the K-Player game this allowed us to fix the population distribution as a male majority. In a mixed strategies K-Player game one must also consider that the distribution over actions of the other players is probabilistic. In this case, one could not assume that men *always* have the majority. Thus, the size of B and the rank order over the actions of each individual given the possible distributions of actions of the other individuals would need to be considered.

5 Summary

This paper demonstrated how to form K-Player games by aggregating the utilities or expected utilities of a collective. A collective is formed by partitioning a population according to their preference order over the possible actions given the strategies of the others. The strategy of player K_i in a K-Player game includes the size of collective K_i. We showed that the Nash equilibria for the population can be found using the K-Player game, and that the equilibria are the same as the N player game. Our approach enables one to find the Nash equilibria in a large population efficiently, without constraining individuals to share strategies with only their neighbors.K-Player games will be most appropriately applied in situations in which each player is only interested in the portion of the population that is taking on each action.

References

[Arrow, 1950] Arrow, K. J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58.

[Arrow, 1963] Arrow, K. J. (1963). *Social Choice and Individual Values, Second edition.* John Wiley & Sons, 2nd edition.

[Daskalakis et al., 2006] Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2006). The complexity of computing a nash equilibrium. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 71–78, New York, NY, USA. ACM.

[Feldman and Serrano, 2007] Feldman, A. M. and Serrano, R. (2007). Arrowś impossibility theorem: Two simple single-profile versions. Working Papers 2007-07, Instituto Madrileo de Estudios Avanzados (IMDEA) Ciencias Sociales.

[Greene, 2010] Greene, K. (2010). *Collective belief models for representing consensus and divergence in communities of Bayesian decision-makers*. PhD thesis, University of New Mexico Department of Computer Science.

[Greene et al., 2010] Greene, K., Kniss, J., and Luger, G. (2010). Representing diversity in communities of bayesian decision-makers. In *Proceedings of the IEEE Conference on Social Computing (SocialCom-2010), Social Intelligence and Networking Symposium*.

[Howard and Matheson, 1984] Howard, R. and Matheson, J., editors (1984). *Readings on the Principles and Applications of Decision Analysis*, volume 2, pages 6–16. Strategic Decisions Group.

[Hylland and Zeckhauser, 1979] Hylland, A. and Zeckhauser, R. J. (1979). The impossibility of bayesian group decision making with separate aggregation of beliefs and values. *Econometrica*, 47(6):1321–36.

[Kearns et al., 2001] Kearns, M. J., Littman, M. L., and Singh, S. P. (2001). Graphical models for game theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, UAI '01, pages 253–260, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

[Leyton-Brown and Shoham, 2008] Leyton-Brown, K. and Shoham, Y. (2008). *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan and Claypool Publishers.

[Matzkevich and Abramson, 1992] Matzkevich, I. and Abramson, B. (1992). The topological fusion of bayes nets. In *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*.

[Pearl, 1988] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman.

[Pennock and Wellman, 1999] Pennock, D. M. and Wellman, M. P. (1999). Graphical representations of consensus belief. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*.

[Pennock and Wellman, 2005] Pennock, D. M. and Wellman, M. P. (2005). Graphical models for groups: belief aggregation and risk sharing. *Decision Analysis*, 2(3):148–164.

[Seidenfeld et al., 1989] Seidenfeld, T., Kadane, J., and Schervish, M. (1989). On the shared preferences of two bayesian decision makers. *The Journal of Philosophy*, 86(5):225–244.

[Shachter, 1986] Shachter, R. D. (1986). Evaluating influence diagrams. *Operations Research*, 34(6):871–882.