

Bargaining for Coalition Structure Formation

Shaheen Fatima¹ and Tomasz Michalak² and Michael Wooldridge³

Abstract. Many multiagent settings require a collection of agents to partition themselves into coalitions. In such cases, the agents may have conflicting preferences over the possible coalition structures that may form. We investigate a noncooperative bargaining game to allow the agents to resolve such conflicts and partition themselves into *non-overlapping* coalitions. The game has a *finite horizon* and is played over discrete time periods. The bargaining agenda is defined *exogenously*. An important element of the game is a parameter $0 \leq \delta \leq 1$ that represents the probability that bargaining ends in a given round. Thus, δ is a measure of the degree of democracy (ranging from *democracy* for $\delta = 0$, through increasing levels of *authoritarianism* as δ approaches 1, to *dictatorship* for $\delta = 1$). For this game, we focus on the question of how a player's position on the agenda affects his *power*. We also analyse the relation between the distribution of the power of individual players, the level of democracy, and the welfare efficiency of the game. Surprisingly, we find that purely democratic games are welfare inefficient due to an uneven distribution of power among the individual players. Interestingly, introducing a degree of authoritarianism into the game makes the distribution of power more equitable and maximizes welfare.

1 Introduction

In this paper, we focus on the problem of how a group of agents can partition themselves into a coalition structure. Specifically, we are interested in the the following scenario. There is a set N of agents who want to decide how to partition themselves into *non-overlapping coalitions*. There are *externalities* in that each agent has preferences not just over coalitions but over the possible coalition structures, i.e., the partitions of N . Conflicts arise because the agents' preferences are not identical. We aim to study how such conflicts can be resolved through a process of noncooperative bargaining. We assume that utility is *not transferable*, i.e, payoffs assigned to one agent cannot be assigned to another agent.

In order to model externalities and non-transferable utilities, we introduce *coalition structure games* (CSGs) that encompass many important classes of coalitional games including hedonic games [4] and NTU (non-transferable utility) games in partition function form [13]. Using a CSG as the underlying game, we investigate the following noncooperative bargaining protocol.

Our multilateral bargaining protocol has a *finite horizon* and is built on Rubinstein's alternating offers protocol [15] for bilateral bargaining. It runs in a series of rounds and the agents take turns to propose an offer, i.e., a coalition structure. The order in which the agents are called to make offers, i.e., the bargaining *agenda*, is defined *exogenously*. An important element of our game is the param-

eter $0 \leq \delta \leq 1$, which represents the probability that bargaining ends in a given round. Thus, δ is a means to control the subset of agents that can make proposals during the bargaining and can be interpreted as a measure of democracy within the game. For $\delta = 1$, only one agent gets the chance to propose—we call these *dictatorship games* in that a single agent is able to influence the outcome. For $\delta = 0$, all the agents get a chance to propose—we call these *democratic games* in that every agent is able to control the outcome. In between, i.e., for $0 < \delta < 1$, we have games that are neither purely dictatorship nor purely democratic but with a degree of *authoritarianism*.

For exogenous agendas, it is well known that the outcome of a bargaining depends on the agenda [3]. Given this, an individual agent will want to know what agenda position is best for him. On the other hand, from the system's perspective, it is necessary to know what agenda will maximize *social welfare*. To these ends, our first objective is to analyze how an agent's position on the agenda influences his *bargaining power* (i.e., his ability to secure a favourable outcome). As an extreme example, for $\delta = 1$, the dictator always secures his most preferred coalition structure while all other agents are completely powerless. However, for any $\delta < 1$, all agents in the agenda possess some degree of power, perhaps, however, not allocated evenly. Thus, our second objective is to analyse the relationship between the distribution of power, the level of democracy (embodied in δ), and the *efficiency* of the game (i.e., how far the bargained structure is from the socially optimal structure).

This paper provides the first quantitative analysis of power and efficiency in the context of noncooperative bargaining for coalition structure formation. To date, the only such analysis of a mechanism for bargaining coalition structure that we are aware of was done by Bloch and Rottier [3] for simple coalitional games, where coalitions can only have binary values. Many multi-agent systems, however, cannot be modeled as simple games. Thus, we consider bargaining in the context of coalition structure games.

Our most important results, from both theoretical analysis and simulations, can be summarised as follows:

- Assuming purely democratic games ($\delta = 0$), the first mover has the lowest power and power increases monotonically with the position on the agenda. Surprisingly, the maximum is reached not for the last mover, n , but for the two last movers, $n - 1$ and n .
- An even more surprising result is that an agent's average power (average taken over all possible player preference combinations) increases with the number of agents in the game. It is always profitable for incumbents to invite more agents to the game as long as the newcomers will be given the very first positions on the agenda.
- For $\delta \approx 0$, power monotonically increases from the first to the last player on the agenda. As δ increases, the difference between the first and last mover's powers decreases and power becomes more balanced. With further increase in δ , power monotonically decreases from the first to the last player. In other words, as the

¹ University of Loughborough, UK, s.s.fatima@lboro.ac.uk

² University of Oxford, UK, email: tomasz.michalak@cs.ox.ac.uk

³ University of Oxford, UK, email: mjw@cs.ox.ac.uk

degree of authoritarianism increases (i.e., as δ approaches 1), there is a shift in power from the last movers toward the first movers.

- The above result is related to a notable finding that concerns *efficiency*. While it is widely accepted in economics and political science [17] that democracy is, in principle, efficient, this is not so in our model. Surprisingly, it is authoritarianism that maximizes expected system welfare. We show that the reason behind this result is to be found in the distribution of power. Democracy becomes inefficient if power is unevenly distributed, and some level of authoritarianism is needed to balance power and maximize welfare.

2 The Model

We begin by defining CSGs and the bargaining game (BG).

Coalition Structure Game (CSG) A coalition structure π is an exhaustive partition of the set of players $N = \{1, \dots, n\}$ into disjoint coalitions; let $\Pi(N)$ denote the set of coalition structures over N and $|\Pi(N)| = \text{Bell}(n) = e^{-1} \sum_{k=0}^{\infty} k^n (k!)^{-1}$ the number of possible coalition structures.

Definition 1 A CSG is a tuple $\mathbb{G} = \langle N, \succ_1, \dots, \succ_n \rangle$ where $N = \{1, \dots, n\}$ is the set of players, and $\succ_i \subseteq \Pi(N) \times \Pi(N)$ is a complete, non reflexive, and transitive preference relation for player $i \in N$, with the interpretation that if $\pi_1 \succ_i \pi_2$, then player i prefers coalition structure π_1 more than the structure π_2 .

A player's preference for a coalition structure is given by a unique integer rank. A player's most preferred structure is the one with rank 1, then with rank 2, and so on. Different players will have different ranks (preference orderings) for the possible structures. Let $r_i(\pi)$ denote player i 's rank for a structure π .

2.1 The Bargaining Game

For a CSG, we explore the following non-cooperative bargaining game for forming a coalition structure. This is a finite-horizon game in which the players take turns in proposing offers where an offer is a coalition structure. The sequence in which the players are called to make offers is called the *bargaining agenda*. The bargaining agenda ρ is a permutation of the first n integers.

The game can run for at most n time periods. Bargaining starts in the first time period and proceeds as described in Algorithm 1. To begin, all the players in N are in the set IN . The set OUT is initially empty (Line 1 in Algorithm 1). At $t = 1$, mover 1 offers a coalition structure $\pi \in \Pi(in)$. After an offer is proposed, the game will end with probability δ . With probability $(1 - \delta)$, it will continue to the next round when mover 2 will propose an offer.

In general, whether or not the game will end at a time t is determined by a 'RandomEvent' (Line 4 in Algorithm 1). If the event does not occur, the game continues at Line 7. In the for loop of Line 7, the players $\rho_{t+1}, \dots, \rho_n$ receive the offer. These players can accept or reject the current offer. If all these players accept, then the game ends and the structure $\pi \cup \{\{i\} : i \in out\}$ is the outcome of the game. But if at least one of these players rejects, then the proposing player is moved from the set IN to the set OUT , time is incremented and we go to Line 2. Then, the next player on the agenda, i.e., ρ_{t+1} proposes an offer. This process repeats until all the subsequent players (i.e., $\rho_{t+1}, \dots, \rho_n$) accept a proposal, or the time $t = n$ is reached. If we reach $t = n$, the game ends and the structure $\pi \cup \{\{i\} : i \in OUT\}$ is implemented.

Definition 2 A BG is a 4-tuple $G = \langle N, \succ_1, \dots, \succ_n, \delta, \rho \rangle$.

Algorithm 1 Bargaining Game G

Require: A δ and a given agenda, i.e., an ordering ρ of players

- 1: $OUT \leftarrow \emptyset$; $IN \leftarrow N$; $t \leftarrow 1$;
 - 2: **while** $t \leq n$ **do**
 - 3: Player ρ_t proposes a coalition structure π made of players in IN .
 - 4: **if** RandomEvent occurs (with probability δ) **then**
 - 5: The game ends. Go to Line 14.
 - 6: **end if**
 - 7: **for** $i = t + 1$ to n **do**
 - 8: **if** (Player ρ_i rejects the proposal) **then**
 - 9: $OUT \leftarrow OUT \cup \{\rho_i\}$; $IN \leftarrow IN - \{\rho_i\}$; $t \leftarrow t + 1$
 - 10: Go to Line 2.
 - 11: **end if**
 - 12: **end for**
 - 13: **end while**
 - 14: The game ends and $\pi \cup \{\{i\} : i \in OUT\}$ is implemented.
-

We will denote a BG as $G(n, \delta, \rho)$, $G(n, \rho)$, $G(n)$, $G(\rho)$, or just G when the other parameters are clear from context. P will denote the set of all possible preference combinations for the n players. All other parameters remaining the same, we can obtain different BGs by varying \succ_1, \dots, \succ_n . We have $|P| = ((\text{Bell}(n))!)^n$ ($\text{Bell}(n)$ is the n th Bell number) elements in P , i.e., we can have $|P|$ possible BGs. \mathcal{G} denotes the set of all $|P|$ games.

The Significance of 'RandomEvent': 'RandomEvent' is used to model a wide spectrum of BGs. For $\delta = 1$, only mover 1 is given a chance to propose and then the game ends; no other player has any say in the outcome of the game. This is a dictatorship game. For $\delta = 0$, all the players are given a chance to propose. Here, G is an n time period game. Since every player can exercise control over the outcome of the game, these are democratic games. For $0 < \delta < 1$, the number of time periods in G will vary between 1 and n ; only the first few players on the agenda will get a chance to propose but the rest will not. These are *authoritarian* games.

3 Equilibrium Strategies for G

For a BG $G = \langle N, \succ_1, \dots, \succ_n, \delta, \rho \rangle$, we will show how to obtain equilibrium strategies. Let $\pi_t^*(G)$ denote the equilibrium offer for time $1 \leq t \leq n$. But when G is evident from context, we will simply denote this offer as π_t^* . Let $er_i(\pi_t)$ denote player i 's expected rank from an offer π_t made at time $t < n$. Here, $er_i(\pi_t)$ is defined as:

$$er_i(\pi_t) = \begin{cases} r_i(\pi_t) & \text{if } r_{\rho_k}(\pi_t) \leq r_{\rho_k}(\pi_{t+1}^*) \\ & \text{for } k > t \\ \delta r_i(\pi_t) + (1 - \delta) r_i(\pi_{t+1}^*) & \text{otherwise} \end{cases}$$

For $t = n$, there is only one possible offer: $\{\{i\} : i \in N\}$. Thus, in the last round, for each $i \in N$, $er_i(\{\{i\} : i \in N\}) = r_i(\{\{i\} : i \in N\})$. We are now ready to formalize the equilibrium strategies. In the following text, IN_t (OUT_t) will denote the set of players in IN (OUT) at time t .

Theorem 1 For a G , the following strategies form a subgame perfect Nash equilibrium. At $t = n$, mover n will propose the structure $\pi_n^* = \{\{i\} : i \in N\}$ and all other players will accept. At $t < n$, ρ_t will propose the structure $\pi_t^* = \bar{\pi}_t \cup \{\{i\} : i \in OUT_t\}$ where

$$\bar{\pi}_t \in \arg \min_{\pi \in \Pi(IN_t)} er_{\rho_t}(\pi) \text{ s.t. } er_{\rho_i}(\pi) \leq er_{\rho_i}(\pi_{t+1}^*) \text{ for } i > t.$$

Each player ρ_i ($i > t$) will accept an offer π_t made at time t if $er_{\rho_i}(\pi_t) \leq er_{\rho_i}(\pi_{t+1}^*)$, and reject otherwise.

Proof 1 We use backward induction. Consider the last round $t = n$ when there will be only one player (i.e., ρ_n) in IN_n and the first $n - 1$ players will be in OUT_n . The only possible offer ρ_n could propose is $\pi_n^* = \{\{i\} : i \in N\}$. This would be accepted by all other players and outcome would comprise n singleton coalitions.

For each previous round $t < n$, ρ_t will propose a structure that minimizes his expected rank subject to giving each subsequent player ρ_i ($i > t$) an expected rank of at least $er_{\rho_i}(\pi_{t+1}^*)$. Since the players in OUT_t are singleton coalitions, the equilibrium offer π_t^* will be:

$$\pi_t^* = \bar{\pi}_t \cup \{\{i\} : i \in OUT_t\}, \text{ where}$$

$$\bar{\pi}_t \in \arg \min_{\pi \in \Pi(IN_t)} er_{\rho_t}(\pi) \text{ s.t. } er_{\rho_i}(\pi) \leq er_{\rho_i}(\pi_{t+1}^*) \text{ for } i > t.$$

Each player ρ_i ($i > t$) will accept an offer π_t made at t if $er_{\rho_i}(\pi_t) \leq er_{\rho_i}(\pi_{t+1}^*)$, and reject otherwise.

Example 1 illustrates the application of Theorem 1.

Structure	(a)			(b)		
	ρ_1	ρ_2	ρ_3	ρ_1	ρ_2	ρ_3
{1, 2, 3}	1	2	1	1	3	2
{1, 2}{3}	4	4	4	2	2	1
{1, 3}{2}	3	3	3	3	1	3
{1}{2, 3}	5	1	5	4	4	4
{1}{2}{3}	2	5	2	5	5	5

Table 1. The players' rankings for all possible coalition structures.

Example 1 For a 3-player game, the preferences are as given in Table 1(a). Let $\delta = 0$. For the last time period, the equilibrium structure will be {1}{2}{3} which is least preferred by mover 2, and ranked 2 by movers 1 and 3. For $t = 2$, the equilibrium structure will be {1}{2, 3} which is most preferred by ρ_2 . Thus, at $t = 1$, ρ_2 will not agree to any other structure. The equilibrium for $t = 1$ will be {1}{2, 3}. Here, ρ_2 secures his most preferred outcome. Next, consider the preferences in Table 1(b). For $t = n$, the equilibrium structure will be {1}{2}{3} which is least preferred by each of the three players. For $t = 2$, the structure will be {1}{2, 3} which is ranked 4 by each player. Thus, at $t = 1$, ρ_1 will propose his most preferred structure {1, 2, 3}, and ρ_2 and ρ_3 will accept because they both prefer it more than {1}{2, 3}. Here, ρ_1 is able to secure his most preferred outcome.

Let $\delta = 0.2$. First, consider the preferences in Table 1(a). For $t = n$, the structure will be {1}{2}{3} which is least preferred by ρ_2 , and ranked 2 by ρ_1 and ρ_3 . For $t = 2$, ρ_2 's equilibrium offer is the structure {1}{2, 3} since his expected rank from it is $0.2 \times 1 + 0.8 \times 5 = 4.2$ while that from {1}{2}{3} is 5. For $t = 1$, ρ_1 's equilibrium offer is his most preferred structure {1, 2, 3} as this gives him the least expected rank of $0.2 \times 1 + 0.8 \times 5 = 4.2$. Thus, if the random event occurs at $t = 1$, the resulting outcome will be {1, 2, 3}. Otherwise, ρ_2 will reject ρ_1 's offer and propose {1}{2, 3}. If the random event occurs at $t = 2$, the resulting outcome will be {1}{2, 3}. Otherwise, ρ_3 will reject ρ_2 's offer and the game will end at $t = 3$ with {1}{2, 3} as the outcome.

Next, for $\delta = 0.2$, consider the preferences in Table 1(b). For $t = n$, the structure will be {1}{2}{3} which is least preferred by all the players. For $t = 2$, ρ_2 's equilibrium offer {1}{2, 3} will be accepted by movers 2 and 3. For $t = 1$, mover 1's offer will be

{1, 2, 3}. Whether the random event occurs or not, the game will end at $t = 1$ with {1, 2, 3} as the outcome.

Proposition 1 If $\delta = 0$ or $\delta = 1$, a game G will result in an immediate agreement, i.e., at $t = 1$. But for $0 < \delta < 1$, an agreement can occur at any time $1 \leq t \leq n$.

Proof 2 If $\delta = 0$, with certainty, we have n time periods. Here, mover 1's equilibrium offer will be accepted by all the subsequent players. If $\delta = 1$, with certainty, only the first mover makes an offer and the game ends at $t = 1$. Next, consider the 3-player game with preferences given Table 1(a). We already demonstrated in Example 1 that an agreement can occur at $t = 1$, $t = 2$, or $t = 3$.

4 Power and Efficiency

We measure a player's power by considering his ability to secure a preferable equilibrium structure averaged over all possible combinations of players' preferences:

Definition 3 Player ρ_i 's power index over the set of games in \mathcal{G} is:⁴

$$\mathbb{P}_{\rho_i}(\mathcal{G}) = 1 - ((\mathbb{E}(\rho_i) - 1) / (\text{Bell}(n) - 1));$$

where $\mathbb{E}(\rho_i)$ denotes ρ_i 's average expected rank in the equilibrium for the games in \mathcal{G} and defined as follows:

$$\mathbb{E}(\rho_i) = \frac{1}{(\text{Bell}(n)!)^n} \sum_{G \in \mathcal{G}} er_{\rho_i}(\pi_1^*(G)).$$

Example 2 For $N = \{1, 2, 3\}$, we have $|\Pi(N)| = 5$ and the number of all preference orderings is $(5!)^3 = 1, 728, 000$. The players' powers are: $\mathbb{P}_1(\mathcal{G}) = 0.61$, $\mathbb{P}_2(\mathcal{G}) = \mathbb{P}_3(\mathcal{G}) = 0.71$ (see Figure 1).

A dictator can secure his most preferred structure irrespective of the preferences of the other players, i.e., $\mathbb{P}_{dictator}(\mathcal{G}) = 1$. But a powerless player has to accept any structure proposed by others, regardless of his preferences, i.e., $\mathbb{P}_{powerless}(\mathcal{G}) = 1/2$. This is because, his average rank, where average is taken over the equilibria for all possible combinations of player preferences, will have exactly the middle rank between 1 and $\text{Bell}(n)$. Thus, for $\delta = 1$, mover 1 is the dictator while the others are powerless.

We note that our method of measuring a player's power is related to the social choice literature, where one of the research problems is to measure the distance between individual preference orderings and a given social preference ordering [9, 11]. To this end, various distance functions were developed [6, 10, 16]. For instance, if we maximize the Kemeny distance function [10] then we choose the social preference ordering that has the highest number of pairwise agreements with individual preference orderings.⁵ In our case, the equilibrium coalition structure is a unique alternative and not a preference ordering over all alternatives. Given a particular combination of preferences, we measure how a given player perceives the quality of the outcome by the rank of the equilibrium coalition structure in the preference ordering of this player. In other words, we count the number of inconsistencies (or pairwise disagreements) between the equilibrium coalition structure (that is chosen, i.e. ranked first, by the protocol) and all other alternatives in the preference ordering of the player. For instance, if players' preference ordering is

⁴ See Section 2.1 for a definition of \mathcal{G} .

⁵ Intuitively, if two preference orderings rank one alternative over the other then we say that there is pairwise agreement between both orderings.

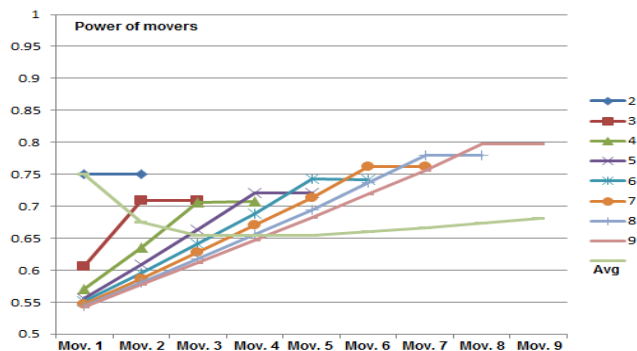


Figure 1. The power in games of different size ($\delta = 0$).

$\pi_1 \succ \pi_2 \succ \pi_3 \succ \pi_4 \succ \pi_5$ and the equilibrium coalition structure is π_4 then we assume that this outcome has the following inconsistencies with the player's preferences: $\pi_4 \succ \pi_1$ instead of $\pi_1 \succ \pi_4$, $\pi_4 \succ \pi_2$ instead of $\pi_2 \succ \pi_4$ and $\pi_4 \succ \pi_3$ instead of $\pi_3 \succ \pi_4$. Naturally, our method of measuring power, similarly to Kemeny [10], Slater [16] or Dodgson [6] distance functions, have various drawbacks. For instance, they all assume that pairwise (dis)agreements between any two alternatives are counted equally, irrespective of how "far" they are in preference orderings [8].

For measuring the quality of an outcome from the systems' perspective, we use *social welfare*. The welfare of a structure is the sum of the individual players' ranks for it. The lower the sum, the higher the welfare. A globally optimal structure π_{SW} maximizes welfare.

Definition 4 A structure is welfare maximizing if it minimizes the sum of the players' ranks, i.e., $\pi_{SW} = \arg \min_{\pi \in \Pi(N)} \sum_1^n r_i(\pi)$.

A bargained structure may not be the same as π_{SW} . In order to measure how far a bargained structure is from π_{SW} , we define *efficiency ratio*.

Definition 5 For a given n and δ , the efficiency ratio, $\mathcal{E}(G)$, is the ratio of the sum of the players' ranks for the globally optimal structure and the sum of ranks for the bargained structure, i.e., we have:

$$\mathcal{E}(G) = \left(\sum_{i=1}^n r_i(\pi_{SW}) \right) / \left(\sum_{i=1}^n er_i(\pi_1^*(G)) \right)$$

Since the sum of ranks for a bargained structure can never be lower than that for the globally optimal structure, we have $\mathcal{E}(G) \leq 1$ for any n and δ . Below, we analyse the power and efficiency of our game starting with the description of the simulation setup.

4.1 Simulation setup

The model was implemented in C++ utilizing the message passing interface (MPI). The computations were run on a cluster in which all computers were equipped with a 4-core AMD Opteron processor (2.0Ghz - 2.3Ghz) and 16 or 32 GB of RAM. At the height of the computations, 72 cores were utilized. Due to the nature of the model, we achieved a very efficient parallelization with an almost ideal linear speed-up. For $n = 3$, it was possible to calculate the whole state space of preferences (that is, $5!^3$ possible combinations for all 3 players, see Example 2). However, already for 4 players, the number of all possible combinations of preferences equals $(15!)^4 = 2.9 \times 10^{48}$.




Figure 2. The average efficiency ratios for games of different size ($\delta = 0$).

Consequently, for $n \geq 4$, we approximate \mathbb{P}_{ρ_i} using Monte Carlo sampling. Each iteration is done as follows. In the first step, we sample random combinations of preference orderings (of all the agents). Then, given these preferences and the assumed values of ρ and δ , we solve the game with backward induction and compute the resulting powers of movers. Next, we update the average and assess the quality of the approximation by comparing the average power of movers $n - 1$ and n .⁶

For $n = 4$, at least 2,000,000 random preferences were sampled using the Knuth shuffle algorithm. For larger n , the number of iterations were increased to ensure convergence. For $n = 9$, we sampled 4,000,000 per case (δ) and the simulation took 24 hours by utilizing 36 cores. The largest n calculated was $n = 14$ and the program required approximately 64 GB of RAM and took over 20 minutes for only a single iteration.

4.2 Democracy ($\delta = 0$)

Power: We can distinguish two types of power. For the preferences of Table 1(a), mover 2 exercises his *power to reject* ρ_1 's proposal. On the other hand, mover 1 has the *power to propose* his best choice which does not have to necessarily correspond to the first choice of mover 2 or the subsequent movers. This, indeed happens for Table 1(b), where mover 1 proposes $\{1, 2, 3\}$ and cannot be rejected since both other players like to cooperate with him, albeit in pairs.

From the definition of the BG, it is easy to observe that, for $n \geq 2$, mover 1 has only the power to propose, while mover n has only the power to reject. The following holds:

Proposition 2 For $n = 2$, mover 1's power is equal to mover 2's power.

Sketch of proof: Consider the number of preference orderings in which both movers can exercise their power. We have only two structures $\pi_1 = \{\{1, 2\}\}$ and $\pi_2 = \{\{1\}, \{2\}\}$ and four possible preference orderings: $(1; 2)_1 = (\pi_1 \succ \pi_2; \pi_1 \succ \pi_2)$, $(1; 2)_2 = (\pi_2 \succ \pi_1; \pi_2 \succ \pi_1)$, $(1; 2)_3 = (\pi_1 \succ \pi_2; \pi_2 \succ \pi_1)$, and $(1; 2)_4 = (\pi_2 \succ \pi_1; \pi_1 \succ \pi_2)$. Here, we have $\mathcal{G} = \{G(N, (1; 2)_1), G(N, (1; 2)_2), G(N, (1; 2)_3), G(N, (1; 2)_4)\}$. There is no conflict for the first two games and for each of these two games, the outcome is a structure that is most preferred (i.e., ranked 1) by both ρ_1 and ρ_2 . For $G(N, (1; 2)_3)$, ρ_2 has the power to reject ρ_1 's offer in order to bring about his most preferred structure

⁶ In Proposition 3 we show that these powers should be the same.




Figure 3. Power and efficiency in G for $n = 3$.

$\{\{1\}, \{2\}\}$. Thus, the outcome of this game will be $\{\{1\}, \{2\}\}$ which is ranked 1 by ρ_2 and 2 by ρ_1 . Here, ρ_1 is powerless. The opposite happens for $G(N, (1; 2)_4)$. Now, ρ_1 will offer his most preferred structure $\{\{1\}, \{2\}\}$ and ρ_2 cannot gain anything by rejecting ρ_1 's offer. In other words, ρ_2 has the power to propose. Thus, the outcome will be $\{\{1\}, \{2\}\}$ which is ranked 1 by ρ_1 and 2 by ρ_2 . Therefore, $\mathbb{E}(\rho_1) = \mathbb{E}(\rho_2) = 5/4$ and $\mathbb{P}_{\rho_1}(G) = \mathbb{P}_{\rho_2}(G) = 3/4$. \square

Proposition 3 For $n \geq 2$, the last two players ρ_{n-1} and ρ_n have equal power, i.e., $\mathbb{P}_{\rho_{n-1}}(G) = \mathbb{P}_{\rho_n}(G)$.

Sketch of proof: Recall that, at time $n - 1$, both ρ_{n-1} and ρ_n have preference orderings over $Bell(n)$ coalition structures but, as per the rules of bargaining, only two coalition structures are feasible: $\pi = \{\{1\}, \{2\}, \dots, \{n\}\}$ and $\pi' = \{\{1\}, \{2\}, \dots, \{n-2\}, \{n-1, n\}\}$. Since we consider all possible combinations of preference orderings, the number of times $\pi \succ \pi'$ for ρ_{n-1} will be the same as the number of times $\pi \succ \pi'$ for ρ_n . Thus, $\mathbb{P}_{\rho_{n-1}}(G) = \mathbb{P}_{\rho_n}(G)$. \square

The fact that the ρ_{n-1} and ρ_n have equal power is convenient when computing power in bigger games. Figure 1 presents the results for $n = 2, 3, \dots, 9$, from which we observe the following:

- (a) For all $n \geq 2$, the power of movers n and $n - 1$ is identical (Proposition 3).
- (b) For all $n \geq 3$, the power of movers $1 \dots, n - 2$ increases monotonically, i.e., $\forall_{3 \leq i \leq n-2} \mathbb{P}_{\rho_i}(G) < \mathbb{P}_{\rho_{i+1}}(G)$.
- (c) For all $n \geq 2$, the power of the first mover decreases with n , i.e., $\forall_{n \geq 2} \mathbb{P}_{\rho_1}(G(n)) < \mathbb{P}_{\rho_1}(G(n+1))$.
- (d) For all $n \geq 4$, the power of $n - 1$ last movers in the game $G(n)$ increases monotonically w.r.t. the power of $n - 1$ movers in $G(n - 1)$, i.e., $\forall_{1 \leq i \leq n-1} \mathbb{P}_{\rho_{n-i}}(G(n-1)) < \mathbb{P}_{\rho_{n-i}}(G(n))$.
- (e) For all $n \geq 4$, the power of an average agent in the game increases with n , i.e., $\forall_{n \geq 4} \frac{\sum_{i=1}^n \mathbb{P}_{\rho_i}(G)}{n} < \frac{\sum_{i=1}^{n+1} \mathbb{P}_{\rho_i}(G)}{n+1}$.

Observation (e) is especially interesting. Counter-intuitively, the more the players in the game, the more powerful, on average, they become. Should the pattern in Figure 1 hold for $n \geq 10$, we conjecture that, with n going to infinity, the power of the first mover converges to a number close to 0.5, that is, this mover gradually becomes powerless. At the same time, the power of the last two movers converges to a number close to 1, that is, their power becomes close to the dictatorship. In addition, since the power of all the remaining movers (from mover 2 to mover $n - 2$) increases monotonically, it can be expected that the average power in the system, with n going to infinity, converges to ≈ 0.75 .




Figure 4. Power and efficiency in G for $n = 7$ and $0 \leq \delta \leq 1$.

This analysis leads to interesting strategic conclusions: if new players are to be added to the game, then, from the point of view of an incumbent player, it is best to add them at the beginning of the agenda. Furthermore, it is best for the incumbents to enlarge the game as much as possible.

Let us now study the relationship between power and efficiency.

Efficiency: Figure 2 shows the average efficiency ratios for games of $n = 2, \dots, 9$ players. Here, the efficiency ratio decreases with n . The reasons can be sought in the distribution of power in the game. As shown in Figure 1, the discrepancies between the power of agents increase with n . This means that, more and more often, powerful agents are able to secure favourable outcomes at the expense of powerless agents—a conflict which results in the overall efficiency loss.

Consider the example in Table 1(a). Here, mover 2 is able to secure his most preferred structure $\{\{1\}\{2, 3\}\}$, but this structure is ranked 5 by both movers 1 and 3. This yields a welfare of $5 + 1 + 5 = 11$. Should mover 2 have less power than mover 1, the grand coalition might be the outcome with a resulting welfare of $1 + 2 + 1 = 4$.

Next, we consider $\delta \geq 0$ as a potential method to balance power.

4.3 Authoritarianism ($\delta > 0$)

We analyse the authoritarian games ($\delta > 0$), including its most extreme case—the dictatorship ($\delta = 1$).

Power: Consider the 3-player game for which we computed the exact power. In Figure 3, we plot the power of all the three movers for different values of δ . With growing probability of the random event, mover 1 becomes increasingly powerful while movers 2 and 3 less powerful. The intuition for this has already transpired in Example 1: for $\delta = 0.2$ in Table 1(a), mover 1 knows that mover 2 opts for the very disliked $\{\{1\}\{2, 3\}\}$, he is going to propose the most preferred grand coalition with the hope that the random event happens (and the grand coalition results). Interpreting the random event from the

power perspective, a higher δ increases mover i 's ($1 \leq i < n$) power to propose, and decreases mover j 's ($i < j \leq n$) power to reject.

Analysing Figure 3 further, for $\delta \approx 0.2$, the power of mover 2 surpasses the powers of the other players. Finally, as expected, for $\delta = 1$, mover 1 becomes the dictator while movers 2 and 3 become completely powerless. Note that, the power of movers 2 and 3 does not decrease uniformly. For $\delta > 0$, mover 2 becomes more powerful than mover 3. This result can be generalized as follows:

Proposition 4 For $\delta > 0$, the power of mover $n-1$ is always higher than that of mover n , i.e., $\mathbb{P}_{\rho_{n-1}} > \mathbb{P}_{\rho_n}$.

Proof 3 As we saw in the proof of Proposition 3, if mover $n-1$ prefers not to cooperate with mover n , the structure $\{\{1\}\{2\} \dots \{n\}\}$ will be proposed and implemented, irrespective of the preferences of n and whether or not the random event occurs. While the same holds for mover n against mover $n-1$ for $\delta = 0$, this is no longer the case for $\delta > 0$. Now, whenever mover $n-1$ likes cooperation with n while mover n does not, mover $n-1$ is going to propose $\{\{1\}\{2\} \dots \{n-1, n\}\}$ with a hope that the random event happens and this coalition structure becomes implemented. Therefore, the expected rank of mover $n-1$ is higher than the expected rank of mover n . \square

We now analyse how increasing δ influences the distribution of power among the players in a game of $n = 7$ players. This distribution is shown in Figure 4. Initially, for small δ (such as 0.1 and 0.2), the first mover has the lowest power among all and the power grows monotonically as we move down the agenda to subsequent movers in ρ . However, the differences between the powers of movers become smaller and smaller. Interestingly, Proposition 4 implies that mover $n-1$ now becomes the most powerful player in the game. As δ increases, there comes a turning point when the first mover becomes most powerful. For instance, at $\delta = 0.4$, the first mover is the most powerful while mover n is the least powerful. Finally, as in the game of $n = 3$ players, when $\delta = 1$, the first mover becomes the dictator and all the other movers are completely powerless.

Efficiency: We suggested in Section 4.2 that the decreasing efficiency with increasing n can be caused by increasing discrepancy between the powers of movers. The results in Figure 4 confirm this conjecture for $n = 7$. In particular, the upper part of this figure shows the distribution of power for all seven movers in the game, given different values of δ . The lower part of Figure 4 shows the level of welfare corresponding to each value of δ . The optimal efficiency is reached for a δ for which the standard deviation between the powers of agents in the game is the smallest.

5 Related Work

A lot of the literature on hedonic games has focused on analysing various stability concepts for coalition structures [4, 1]. This research shows conditions for the existence of stable structures and studies their properties. However, the actual procedure of how such stable structures emerge is usually left out of the analysis. In contrast, this topic is the focal point of literature on games for coalition structure formation, to which our paper contributes.

[5] studied the efficiency of equilibrium coalition structures in the spirit of Ray and Vohra's [14] equilibrium binding agreements and von Neumann and Morgenstern's stable set. [5] shares with our model the assumption of players' farsightedness. However, the coalition structure formation process in [5] is unstructured, i.e., bargaining is not conducted through a protocol. Also, they deal with a special class partition function games while we deal with CSGs.

To measure how far a bargained outcome will be from a social optimum, Koutsoupias and Papadimitriou [12] defined *price of anarchy* and measured it in terms of the ratio of the worst possible Nash equilibrium and a social optimum. They derived upper and lower bounds for this ratio for resource sharing in *network games*.

For coalitional resource games (a form of NTU games), [7] analyzed a protocol for bargaining cooperation structures. Our protocol is similar to this in terms of the rules of bargaining. However, in [7], $\delta = 0$, but in our protocol $0 \leq \delta \leq 1$. [7] show that the negotiation outcome satisfies desirable properties: Pareto optimality, dummy player, and pseudo-symmetry. In contrast, our focus is on the relation between agenda, power, and efficiency.

6 Conclusions

We explored a noncooperative game for bargaining a coalition structure. This is a finite horizon alternating offers game with an exogenous agenda. In each round, the game ends with a certain probability. We obtained equilibrium for the game, showed how a player's position on the agenda affects his power, and analysed the relationship between the distribution of the power of individual players, the level of democracy, and the welfare efficiency of the protocol. We found that purely democratic games are welfare inefficient. Introducing a degree of authoritarianism into the protocol makes the distribution of power more equitable and maximizes welfare.

While our work is primarily theoretical in nature, we believe that studies on the relationship between agenda and power may shed light on some real-world phenomena. In particular, the approach developed in our paper could be useful, for instance, to get some insight into the distribution of power in structured multilateral negotiations on oligopolistic markets (such as armaments), where the government aims to distribute contracts among consortia of companies.

ACKNOWLEDGEMENTS

Tomasz Michalak & Michael Wooldridge were supported by the European Research Council under Advanced Grant 291528 ("RACE").

REFERENCES

- [1] H. Aziz and F. Brandt, Existence of stability in hedonic coalition formation games, *AAMAS*, 763–770, 2012.
- [2] H. Aziz and F. Brandt and H. Seedig, Optimal Partitions in Additively Separable Hedonic Games, *CoRR*, abs/1005.4540, 2010.
- [3] F. Bloch and S. Rottier, Agenda control in coalition formation, *Social Choice Welfare*, 19:769–788, 2002.
- [4] A. Bogomolnaia and M. Jackson, The stability of hedonic coalition structures, *GEB*, 38:201–230, 2002.
- [5] E. Diamantoudi and L. Xue, Coalitions, agreements and efficiency, *Journal of Economic Theory*, 136(1):105–125, 2007.
- [6] Dodgson, S., A method of taking votes on more than two issues. In: Black, D. *The theory of committees and elections*, CUP, 1958.
- [7] P. Dunne and S. Kraus and E. Manisterski and M. Wooldridge, Solving Coalitional Resource Games, *AI journal*, 174:20–50, 2010.
- [8] Farnoud, F., O. Milenkovic, and B. Touri, A Novel Distance-Based Approach to Constrained Rank Aggregation, <http://arxiv.org/abs/1203.6371>, 2012.
- [9] Fishburn, P., Condorcet social choice functions, *SIAM J Appl Math* 33:469–489, 1977.
- [10] Kemeny, J., Mathematics without numbers, *Daedalus* 88:577–591, 1959.
- [11] Klamler, Ch., The Dodgson ranking and its relation to Kemeny's method and Slaters rule, *Social Choice and Welfare* 23(1):91–102, 2004.
- [12] E. Koutsoupias and C. Papadimitriou, Worst-case equilibria, *STACS-99*, 404–413, 1999.
- [13] D. Ray, A Game-Theoretic Perspective on Coalition Formation, Oxford University Press, 2007.
- [14] D. Ray and R. Vohra, A Theory of Endogenous Coalition Structures, *Games and Economic Behavior*, 26(2): 286–336, 1999.
- [15] A. Rubinstein, Perfect Equilibrium in a Bargaining Model, *Econometrica*, 50:97–109, 1982.
- [16] P. Slater, Inconsistencies in a schedule of paired comparisons, *Biometrika* 48:303–312, 1961.
- [17] D. Wittman, Efficiency of Democracy, Eds. C. Rowley and F. Schneider, *The Encyclopedia of Public Choice*, Springer US, 510–513, 2003.