

On the stability of an Optimal Coalition Structure

Stéphane Airiau¹ and Sandip Sen²

Abstract. The two main questions in coalition games are 1) what coalitions should form and 2) how to distribute the value of each coalition between its members. When a game is not superadditive, other coalition structures (CSs) may be more attractive than the grand coalition. For example, if the agents care about the total payoff generated by the entire society, CSs that maximize utilitarian social welfare are of interest. The search for such optimal CSs has been a very active area of research. Stability concepts have been defined for games with coalition structure, under the assumption that the agents agree first on a CS, and then the members of each coalition decide on how to share the value of their coalition. An agent can refer to the values of coalitions with agents outside of its current coalition to argue for a larger share of the coalition payoff. To use this approach, one can find the CS s^* with optimal value and use one of these stability concepts for the game with s^* . However, it may not be fair for some agents to form s^* , e.g., for those that form a singleton coalition and cannot benefit from collaboration with other agents. We explore the possibility of allowing side-payments across coalitions to improve the stability of an optimal CS. We adapt existing stability concepts and prove that some of them are non-empty under our proposed scheme.

1 Introduction

Forming coalitions is an effective means for agents to cooperate: in a coalition, agents may share resources, knowledge, or simply join forces to improve their performance. Given such incentives for coalition formation, the two pertinent questions are 1) what coalitions will form in the society and 2) how to distribute the worth of a coalition? Stability is a key criteria to answer both questions. Another key question is the social objective, e.g. maximizing utilitarian social welfare.

For superadditive games, forming the grand coalition maximizes utilitarian social welfare. Asking whether the grand coalition is stable amounts to asking whether the core [3] is non-empty. When the core is empty, some other stability concept is needed. For example, an element of the kernel, or the nucleolus of the game, which are known to always be non-empty, can be used. Also, the Shapley value can be used, and though it does not provide stability, it is based on a certain measure of fairness.

For non-superadditive games, the grand coalition may not be stable. Instead, several coalitions may coexist in the population, forming a coalition structure (CS). From the point of view of a system designer, it is desirable to form an optimal CS, i.e., one that maximizes social welfare. The search for such optimal coalition structures has been an active area of research in recent years [8]. For that line of research, however, the stability problem is avoided, the agents are assumed to be fully cooperative and as such, they have no incentive

to deviate from a coalition structure that optimizes the performance of the entire society.

Aumann and Drèze studied the stability of games with CS. They assume that the agents agree on a CS first and then solve the issue of the payoff distribution. For instance, the agents may first locate a CS that maximizes utilitarian social welfare and then find a stable payoff distribution for that CS. Another assumption is that members of a coalition share the value of their own coalition, and they can use agents outside of their coalition to negotiate their payoff, i.e., they may refer to opportunities they have outside of their coalition to justify a side payment. Aumann and Drèze showed that one necessary condition for non-emptiness of the core is that the CS formed is the optimal CS. But this condition is not necessary, and the core of a game with CS may be empty. Again, one can use other stability concepts such as the kernel, the nucleolus or the Shapley value to agree on a payoff distribution.

We are interested in non-superadditive games where the agents are forming an optimal CS, and our goal is to have a payoff distribution that combines elements of stability and fairness. We will allow side-payments between members of two different coalitions so that all agents in the population are in equilibrium. We believe that this will increase fairness as it will eliminate the effect of the structure on agent payoffs. For example, it may be possible that in the optimal CS, some agents are forming singletons, or are members of a small coalition. As such, they may not benefit from the cooperation of other agents, even when they have high marginal contributions to a large number of coalitions. However, for the greater good of the population, i.e., maximizing social welfare, they may be forced to stay in that optimal CS. Aumann and Drèze suggested that one reason to consider games with CS is to take into account externalities that are not present in the valuation function. For example, academics may prefer to stay in their own country, hence the CS of academicians may represent coalition of academics by location. They may refer to potential opportunities of working abroad to negotiate their salary. However, if their country does not provide much budget, their salary may be low. In this paper, if the researchers stay in their country, not because of personal preference, but because they believe it will bring a higher value for the entire community, we would like the community to provide them additional funds, hence allowing payments across coalitions. We note that this is not the usual assumption in game theory. However, we believe that in multiagent systems, allowing such transfer of utility across coalitions may be a possible solution to guarantee stability of an optimal CS.

We present a modification of the kernel stability criterion. We will introduce background research in Section 2. Then we will present our stability concept in Section 3, where we providing the definition and some properties, and then we will describe an algorithm to compute a stable payoff distribution. The proof that a set of stable payoff distributions is non-empty relies on this algorithm. We end that section

¹ ILLC - University of Amsterdam, email: s.airiau@uva.nl

² University of Tulsa, email: sandip@utulsa.edu

by suggesting that we could use the stability concept for games with externalities. Finally, Section 4 presents conclusions.

2 Background

We consider games with transferable utility, also called TU games. These games assume that interpersonal comparison of utility and transfer of utility between agents are possible. A TU game is a pair (N, v) where N is the set of n agents, and $v : 2^N \rightarrow \mathbb{R}$ is the worth, or value of a coalition. A coalition structure (CS) is a partition of the agents into coalitions, i.e., a CS $\mathcal{S} = \{C_1, \dots, C_k\}$ where $\forall l \in \{1, \dots, k\} C_l \subseteq N$, $\cup_{l=1}^k C_l = N$, and $(i \neq j) \Rightarrow C_i \cap C_j = \emptyset$. We denote by \mathcal{S}_C the set of all partitions of $C \subseteq N$. In particular, \mathcal{S}_N is the set of all CSs. We denote by $x \in \mathbb{R}^n$ a payoff distribution, the payoff for agent $i \in N$ is x_i and we use the notation $x(C) = \sum_{i \in C} x_i$ for $C \subseteq N$.

A TU game with CS is a triplet (N, v, \mathcal{S}) where (N, v) is a TU game, and $\mathcal{S} \in \mathcal{S}_N$. Such games have been introduced by Aumann and Drèze in [1]. The traditional assumption is that agents first agree on the CS \mathcal{S} , or \mathcal{S} exists due to some externalities not contained in the valuation function. Then, members of each coalition $C \in \mathcal{S}$ negotiate the distribution of the worth of their coalition C . The agents do not threaten to change coalitions during this negotiation, but they try to negotiate a greater share of the valuation of their coalition. Hence, for each coalition in \mathcal{S} , the sum of the members' payoff cannot exceed the value of the coalition. The set of *feasible* payoff vectors is then defined as $X_{(N, v, \mathcal{S})} = \{x \in \mathbb{R}^n \mid \forall C \in \mathcal{S} x(C) \leq v(C)\}$. The *core* of a game with CS (N, v, \mathcal{S}) is defined by $\{x \in X_{(N, v, \mathcal{S})} \mid \forall C \subseteq N, x(C) \geq v(C)\}$. Aumann and Drèze made a link from a game with CS to a special superadditive game (N, \hat{v}) called the superadditive cover [1]. The valuation function \hat{v} is defined as $\hat{v}(C) = \max_{\mathcal{P} \in \mathcal{S}_C} \{\sum_{T \in \mathcal{P}} v(T)\}$ for all coalitions $C \subseteq N \setminus \emptyset$, and $\hat{v}(\emptyset) = 0$. In other words, $\hat{v}(C)$ is the maximal value that can be generated by any partition of C . They showed that $\text{Core}(N, v, \mathcal{S}) \neq \emptyset$ iff $\text{Core}(N, \hat{v}) \neq \emptyset \wedge \hat{v}(N) = \sum_{C \in \mathcal{S}} v(C)$ and that when $\text{Core}(N, v, \mathcal{S}) \neq \emptyset$, then $\text{Core}(N, v, \mathcal{S}) = \text{Core}(N, \hat{v})$. This means that a necessary condition for (N, v, \mathcal{S}) to have a non empty core is that \mathcal{S} is an optimal CS. It is for the games with an empty core that this paper has an interest: as agents cannot find an agreement, they need to relax their requirements for stability. One way to ensure stability is to relax the constraint of the core by allowing payoff distribution such that $\forall C \subseteq N, x(C) \geq v(C) + \epsilon$ holds. This is the idea of the ϵ -core. Another idea by Bachrach et al. is to add some external payments to stabilize the CS [2]. Another possibility is to use a different stability concept, and in this paper, we consider the kernel [3].

We now formally define the concept of the kernel. The **excess** related to coalition C for a payoff distribution x is defined as $e(C, x) = v(C) - x(C)$. We can interpret a positive excess ($e(C, x) \geq 0$) as the amount of *dissatisfaction* or *complaint* of the members of C from the allocation x . For two agents k and l , the **maximum surplus** $s_{k,l}(x)$ of agent k over agent l with respect to payoff distribution x is $s_{k,l}(x) = \max_{C \subseteq N \mid k \in C, l \notin C} e(C, x)$. This maximum surplus can be used by agent k as a measure for its strength over agent l : assuming it is positive and that the agent can claim all of it, agent k can argue that it will be better off in a coalition that does not contain agent l and hence should be compensated with more utility for staying in the current coalition. Two agents k and l that are in the same coalition are in equilibrium when we have either $s_{kl}(x) \geq s_{lk}(x)$ or $x_k = v(\{k\})$. A payoff distribution is in the **kernel** of the game (N, v, \mathcal{S}) when all agent pairs belonging to the same coalition C in

\mathcal{S} are in equilibrium. Note that a payoff distribution that is kernel-stable for the game (N, v, \mathcal{S}) may not be stable for a different game (N, v, \mathcal{S}') with $\mathcal{S} \neq \mathcal{S}'$. Although they use the same valuation function v to argue, the set of agents that are in equilibrium is different, which will have an impact on the payoff distribution. An approximation of the kernel is the ϵ -kernel where the inequality above is replaced by $s_{k,l}(x) + \epsilon \geq s_{l,k}(x)$, where ϵ is a small positive constant. The kernel is always non-empty, it contains the nucleolus, and is included in the bargaining set [6].

One method for computing a kernel-stable payoff distribution is the Stearns method [12]. The idea is to build a sequence of side-payments between agents to decrease the difference of surpluses. At each step of the sequence, the two agents with the largest difference of maximum surplus exchange utility: the agent with smaller maximum surplus makes a payment to the other agent, which decreases their surplus difference. After each side-payment, the maximum surplus over all agents decreases. In the limit, the process converges to an element in the kernel, which may require an infinite number of steps as the side payments may become arbitrarily small. The use of the ϵ -kernel can alleviate this issue. A criteria to terminate Stearns method is proposed in [11], and we present the corresponding algorithm in Algorithm 1.

Algorithm 1: Transfer scheme converging to a ϵ -kernel-stable payoff distribution for the game with CS (N, v, \mathcal{S})

```

compute- $\epsilon$ -kernel( $N, v, \mathcal{S}, \epsilon, x$ )
repeat
  // compute the maximum surplus
  for each coalition  $C \in \mathcal{S}$  do
    for each pair of members  $(i, j) \in C, i \neq j$  do
       $s_{ij}(x) \leftarrow \max_{R \subseteq N \mid (i \in R, j \notin R)} v(R) - x(R)$ 
     $\delta \leftarrow \max_{(i,j) \in C^2, C \in \mathcal{S}} |s_{ij}(x) - s_{ji}(x)|$ ;
     $(i^*, j^*) \leftarrow \mathop{\text{argmax}}_{(i,j) \in N^2} s_{ij}(x) - s_{ji}(x)$ ;
    if  $(x_{j^*} - v(\{j^*\}) < \frac{\delta}{2})$  then
       $d \leftarrow x_{j^*} - v(\{j^*\})$ ;
    else
       $d \leftarrow \frac{\delta}{2}$ ;
    //  $d$  ensures individually rational payoffs
     $x_{i^*} \leftarrow x_{i^*} + d$ ;
     $x_{j^*} \leftarrow x_{j^*} - d$ ;
  until  $\frac{\delta}{v(\mathcal{S})} \leq \epsilon$ ;

```

Algorithm 1 is of exponential complexity since all coalition values need to be checked for computing the maximum surpluses. Note that when a side-payment is made, it is necessary to recompute the maximum surpluses. The derivation of the complexity of the Stearns method to compute a payoff in the ϵ -kernel can be found in [4, 11], and the complexity for one side-payment is $O(n \cdot 2^n)$. Of course, the number of side-payments depends on the precision ϵ and on the initial payoff distribution. Converging to an element of the ϵ -kernel requires $n \log_2(\frac{\delta_0}{\epsilon \cdot v(\mathcal{S})})$, where δ_0 is the maximum surplus difference in the initial payoff distribution. To derive a polynomial algorithm, the number of coalitions must be bounded. The solution used in [4, 11] is to only consider coalitions whose size is bounded in the interval K_1, K_2 . The complexity of the truncated algorithm is $O(n^{2+K_2})$.

The coalition formation protocol proposed by Shehory and Kraus [11] combines a distributed search to find a Pareto optimal CS with kernel-stable payoff distribution. The use of kernel-stable payoff ensures stability. The use of Pareto optimality ensures that changing CS cannot benefit all the agents at the same time. This approach

cannot scale up to large numbers of agents due to the complexity. By bounding the coalition size in an interval $[K_1, K_2]$, the computation of the kernel can be reduced to polynomial time in the number of agents [4], though the order of this polynomial, K_2 , can be high. One sticky issue is that the CS formed may not be optimal. In our work, we will assume that the agents first find a CS with optimal value, using for example the algorithm by Rahwan et al. [8].

3 A new stability criterion

We will now motivate the need for extending the concept of the kernel payoff distribution and then introduce the necessary terminology.

3.1 Motivations for an extension

We propose the following desirable properties of a stable payoff distribution for a coalition structure:

Efficiency: The payoff distribution should correspond to an efficient agent society.

Global equilibrium: All agents in the population, and not just the agents belonging to the same coalition in the current CS, must be in equilibrium.

Value function fairness: The payoff distribution should reflect the overall global properties of the valuation function and not just the valuations of the coalitions present in any one particular CS.

As the kernel computation takes any CS as input, it has no control over societal efficiency. The kernel satisfies value function fairness as the maximum surplus is a maximization over a set of coalitions, hence, in that respect, the kernel satisfies the value function fairness. However, the size of the set of coalitions analyzed depends on the coalition currently formed by the agents: the more members in the coalition, the larger the set of coalitions to be analyzed. If the CS formed is the CS containing the grand coalition, all the other coalitions are taken into account to compute the maximum surplus. If the agent is forming a singleton coalition, however, there is no need to compute a maximum surplus. The kernel does not satisfy, however, the global equilibrium property.

To provide a more concrete definition of global equilibrium, i.e., to extend the concept of equilibrium from two members of the same coalition to an equilibrium between two agents in the population, we may try to use the maximum surplus. This idea, however, does not work. Given a CS and a payoff distribution u , let us consider two agents i and j that have different maximum surplus, e.g., $s_{ij}(u) > s_{ji}(u)$. If agents i and j are in different coalitions, agent i cannot use this surplus advantage to claim that it would be better off without agent j in a different CS, as its coalition already does not include j !

We believe that the strength of each agent, as in the case of the kernel, should be the value of the excess of a coalition, but that coalition should be chosen differently. To represent a true, resilient equilibrium between two agents in the population, the strength for an agent could be the maximum excess of a coalition containing that agent. More formally, the strength of agent i in the payoff distribution u is $\sigma_i(u) = \max_{C \subseteq N \setminus \{i\}} e(C \cup \{i\}, u)$. If agent i is stronger than agent j , agent i can argue that it deserves more payoff as it is a member of a coalition that can create higher excess. This argument is quite similar to the one for the kernel, except it applies between any two agents including those that are not in the same coalition. When all agents are in equilibrium, no agent wants to change coalition and as a result the CS is stable. More precisely, for all agents to be in equilibrium, the strengths of all agents have to be equal (except for some agent

i such that $u_i = v(\{i\})$). Hence, no agent can claim more payoff and the payoff distribution is stable. As a result, there is equilibrium between all agents in the population, not only between members of each coalition. The kernel⁺ is not dependent on the current CS, but is the property of only the valuation function and the current payoff distribution.

3.2 Definition and Properties

We now formally define our new stability criterion, the kernel⁺. As for the kernel, we start by defining a notion of strength, which will be used to justify a side payment from a weak agent to a strong agent.

Definition 1 (Strength). *The strength of agent i for a payoff distribution u is $\sigma_i(u) = \max_{C \subseteq N \setminus \{i\}} v(C \cup \{i\}) - \sum_{k \in C \cup \{i\}} u_k$*

Note that the strength does not depend on the current CS but is only a property of the valuation function and the current payoff distribution. Computing the strength requires evaluating all coalitions where agent i is present. For a population of n agents, this means that each agent has to evaluate 2^{n-1} coalitions.

Definition 2 (Global Equilibrium). *Given a valuation function v and a payoff distribution u , two agents i and j are in global equilibrium when $\sigma_i(u) \geq \sigma_j(u)$ or $x_i = v(\{i\})$.*

The above definition is the same as the corresponding one for the kernel, except that the maximum surplus $s_{ij}(u)$ between two agents i and j is replaced by the strength $\sigma_i(u)$ of agent i .

Definition 3 (kernel⁺). *For a CS S , a payoff distribution u is in the kernel⁺ when $\sum_{i \in N} u_i = \sum_{C \in S} v(C)$ and any two agents in the population are in equilibrium.*

As the global equilibrium involves all agents in the population, and not only the members of the same coalition, all agents will be sharing the valuation of the CS. In particular, the payoff to a coalition is not necessarily distributed only among its members. Though this is not the usual assumption in traditional coalition formation research, as we are seeking a global measure of stability, it is only fair that the entire worth produced by the CS be shared among all agents in the population according to a global property of the valuation function.

Let us now present an example using the following five-player game $(\{1, 2, 3, 4, 5\}, v)$ with v defined as follows: $v(\{1\}) = 4$, $v(\{j\}) = 0$, $v(\{1, j\}) = 10.5$, $v(\{j, k\}) = 4$, for $(j, k) \in \{2, 3, 4, 5\}$, $j \neq k$, $v(\{2, 3\}) = 8$, $v(\{4, 5\}) = 10$, $v(C) = 11$ for $|C| \geq 3$. The optimal CS is $S = \{\{1\}, \{2, 3\}, \{4, 5\}\}$ with a value of 22. With the kernel, the payoff of agent 1 is $v(\{1\})$ since no side payments across coalitions are possible. Agents 2 and 3 share the value of $v(\{2, 3\})$, similarly agents 4 and 5 share $v(\{4, 5\})$. Because of the symmetric role they play, an equal share will guarantee equilibrium, hence $x = \langle 4, 4, 4, 5, 5 \rangle \in \text{Kernel}(N, v, S)$. In addition, it is not possible to further decrease any excess, hence x is also the nucleolus of the game (N, v, S) . However, this payoff distribution may not be tolerable for player 1: it has the largest value for a singleton coalition, it is present in the coalitions of size 2 that produce the highest value. Agent 1 could be considered at least as good as any other agent, however, two agents have a strictly higher payoff. If the agents were to use the kernel⁺, then agents 4 and 5 should make a side payment to agent 1 since the strength of agent 1 is 2.5 ($e(\{1, 2\}, x) = 10.5 - 4 - 4$), and the strength of agent 4 or 5 is 1.5 ($e(\{1, 4\}, x) = 10.5 - 4 - 5$). This would help agent 1 to accept being in a singleton coalition.

This example shows that the kernel and the nucleolus is not included in the kernel⁺ in general. Of course, for games $(N, v, \{N\})$ (i.e. where the CS is the CS containing the grand coalition), the kernel and the kernel⁺ coincide. The following lemma shows that the core is contained in the kernel⁺.

Lemma 1. $Core(N, v, \mathcal{S}) \subseteq Kernel^+(N, v, \mathcal{S})$

Proof. Let $x \in Core(N, v, \mathcal{S})$, for all $\mathcal{C} \subseteq N$, $x(\mathcal{C}) \geq v(\mathcal{C})$, hence no agent has any positive excess, and the maximum excess is bounded by 0 for each agent. Since the x is in the core, it is a feasible payoff for each coalition in the CS \mathcal{S} , i.e., $\forall \mathcal{C} \in \mathcal{S}, v(\mathcal{C}) = x(\mathcal{C})$. Hence, the maximum excess for each agent is at most 0. It follows that the maximum excess of each agent, i.e., the strength of each agent, is 0 and the agents are in equilibrium. \square

As in the case of the kernel, we can here define an ϵ -kernel⁺ where the global equilibrium is obtained when the difference between the strength of any two agents is less than ϵ .

One possible choice for the CS is an optimal CS, i.e., one that maximizes utilitarian social welfare. By choosing an optimal CS, the agents share the largest payoff they can produce. The values of all CSs are considered when deciding the payoff distribution, and the CS formed then corresponds to the most efficient outcome that the society can generate. The corresponding payoff distribution is then Pareto optimal. Since the entire value of the CS is shared, it is not possible to increase an agent's payoff without decreasing the payoff of at least another agent. Hence, the combination of an optimal CS with an kernel⁺-stable payoff distribution is attractive.

3.3 Algorithm for Computing an kernel⁺-Stable Payoff Distribution

We now present an algorithm that returns an kernel⁺-stable payoff distribution. We start with some arbitrary payoff distribution which is likely to not be at global equilibrium. Because of the similarity with the kernel, the algorithm for computing kernel⁺-stable payoff distribution is based on a sequence of side payments which reduces the difference in strength between pairs of agents as in [12]. As for the kernel, this algorithm requires to read multiple times the entire set of coalitions, which limits its usability. However, the algorithm helps us to prove that the kernel⁺ is non-empty.

First, we present two properties about the dynamics of the strength during a side payment. We will then present the transfer scheme and prove that it returns a kernel⁺-stable payoff distribution. The first property provides bounds of the strength for the two agents that are involved in the side payments before and after such a transaction.

Property 3.1. *After a side payment δ from agent j to agent i with the highest strength, the strength of i strictly decreases and that of j increases.*

Proof. Let δ be a side payment from agent j to agent i with largest strength and let u (respectively w) be the payoff distribution before (respectively after) the side payment: $w_i = u_i + \delta$, $w_j = u_j - \delta$ and $\forall k \neq i, k \neq j, w_k = u_k$. Let $\sigma_l(x)$ denote the strength of agent l for payoff distribution x . Let $\mathcal{C}_l(x)$ be the coalition containing l such that $\sigma_l(x) = e(\mathcal{C}_l, x)$. We now consider two possible situations depending on whether j and i are members of $\mathcal{C}_i(w)$.

1. if $j \in \mathcal{C}_i(w)$: then

$$\sigma_i(w) = v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w)} w_l$$

$$\begin{aligned} &= v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w), l \notin \{i, j\}} w_l + w_i + \delta + w_j - \delta \\ &= v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w)} w_l \end{aligned}$$

By definition, $\sigma_i(u) \geq v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w)} w_l$. If $\sigma_i(u) = v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w)} w_l$, then $\sigma_i(u) \leq \sigma_j(u)$ as $j \in \mathcal{C}_i(w)$. In that case, there should be no side payment from j to i , which would contradict the premise of the proposition. Hence, we must have $\sigma_i(u) > v(\mathcal{C}_i(w)) - \sum_{l \in \mathcal{C}_i(w)} w_l = \sigma_i(w)$.

2. if $j \notin \mathcal{C}_i(w)$:

$$\begin{aligned} \sigma_i(w) &= v(\mathcal{C}_i(w), \mathcal{S}) - \sum_{l \in \mathcal{C}_i(w)} w_l \\ &= v(\mathcal{C}_i(w), \mathcal{S}) - \sum_{l \in \mathcal{C}_i(w)} w_l - \delta \\ &\quad \text{as } i \in \mathcal{C}_i(w) \text{ and } j \notin \mathcal{C}_i(w) \\ &\leq \sigma_i(u) - \delta < \sigma_i(u) \text{ as } \delta > 0 \end{aligned}$$

Hence, $\sigma_i(w) < \sigma_i(u)$, i.e., the strength of i strictly decreases. We will now prove $\sigma_j(w) \geq \sigma_j(u)$.

1. If $i \in \mathcal{C}_j(u)$, then $\sum_{l \in \mathcal{C}_j(u)} w_l = \sum_{l \in \mathcal{C}_j(u)} w_l + \delta - \delta = \sum_{l \in \mathcal{C}_j(u)} w_l$. Hence, $\sigma_j(w) = v(\mathcal{C}_j(u)) - \sum_{l \in \mathcal{C}_j(u)} w_l \leq \sigma_j(u)$.
2. If $i \notin \mathcal{C}_j(u)$. We have by definition of $\sigma_i(w)$: $\sigma_j(w) \geq v(\mathcal{C}_j(u)) - \sum_{l \in \mathcal{C}_j(u)} w_l$. But $\sum_{l \in \mathcal{C}_j(u)} w_l = \sum_{l \in \mathcal{C}_j(u)} w_l - \delta$, and hence, $\sigma_j(w) \geq v(\mathcal{C}_j(u)) - \sum_{l \in \mathcal{C}_j(u)} w_l + \delta$. As $\sigma_j(u) = v(\mathcal{C}_j(u)) - \sum_{l \in \mathcal{C}_j(u)} w_l = \sigma_j(u)$, we have $\sigma_j(w) \geq \sigma_j(u) + \delta$. As $\delta > 0$, we have $\sigma_j(w) + \delta \geq \sigma_j(u)$.

Hence, for both cases, we have $\sigma_j(w) \geq \sigma_j(u)$, i.e. the strength of j increases. \square

This property ensures that a side payment does reduce the difference in strength between two agents, more precisely, it decreases the largest difference in strength. Because of the change in payoff, the value of the strength may have changed for other agents, and we need to ensure that the new largest difference in strength smaller than the old one. The next property ensures that, when agent i receives a payment from agent j , if the strength of any other agent was lower than the one of agent i before the side payment, the strength of that agent in the new payoff distribution remains lower than the one of agent i with the old payoff distribution.

Property 3.2. *Given an initial payoff distribution u and a side payment $\delta = \frac{1}{2}(\sigma_i(u) - \sigma_j(u)) > 0$ from agent i to agent j that produces a new payoff distribution w , for all agents $k \notin \{i, j\}$ such that $\sigma_k(u) \leq \sigma_i(u)$ we have $\sigma_k(w) \leq \sigma_i(u)$.*

Proof. **Case 1:** $((i, j) \in \mathcal{C}_k(w)^2)$ or $(i \notin \mathcal{C}_k(w) \text{ and } j \notin \mathcal{C}_k(w))$

$$\begin{aligned} \sigma_k(w) &= v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} w_l \\ &= v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} w_l \end{aligned}$$

If $v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} w_l > \sigma_i(u)$, then as σ is the strength and $k \in \mathcal{C}_k(u)$, $\sigma_k(u) > \sigma_i(u)$ and then i would not be an agent with the largest strength, which contradicts our hypothesis. Hence $\sigma_k(w) \leq \sigma_i(u)$.

Case 2: $i \in \mathcal{C}_k(w)$ and $j \notin \mathcal{C}_k(w)$ Let us assume $\sigma_i(u) < \sigma_k(w)$.

$$\begin{aligned} \sigma_i(u) &< \sigma_k(w) \\ &< v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} w_l \\ &< v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} u_l - \delta \\ &< \sigma_i(u) - \delta < \sigma_i(u), \text{ a contradiction} \end{aligned}$$

Hence, $\sigma_i(u) \geq \sigma_k(w)$.

Case 3: $i \notin \mathcal{C}_k(w)$ and $j \in \mathcal{C}_k(w)$ Let us assume $\sigma_i(u) < \sigma_k(w)$, then

$$\begin{aligned} \sigma_i(u) &< \sigma_k(w) \\ \sigma_i(u) &< v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} w_l \\ \sigma_i(u) &< v(\mathcal{C}_k(w)) - \sum_{l \in \mathcal{C}_k(w)} u_l + \delta \\ \sigma_i(u) &< \sigma_j(u) + \delta \text{ as } j \in \mathcal{C}_k(w) \\ \sigma_i(u) - \sigma_j(u) &< \delta \\ 2 \cdot \delta &< \delta, \text{ a contradiction} \end{aligned}$$

Hence, $\sigma_i(u) \geq \sigma_k(w)$.

Finally, from cases 1, 2, and 3, we have showed that $\sigma_i(u) \geq \sigma_k(w)$. \square

Then, if the side payment is received by the agent with the highest strength, the first property ensures that the strength of this agent strictly decreases, and the second property ensures that after the side payment, the value of the highest strength cannot increase. Hence, the value of the highest strength is strictly decreasing. Because the value of the strength is bounded by 0, we are guaranteed that a sequence of side payments between the agents with the largest difference in strength will converge to a payoff distribution where agents have equal strength in the limit. This is the idea used by the transfer scheme in Algorithm 2 to compute a payoff distribution in the kernel⁺.

Algorithm 2: Transfer scheme for converging to an ϵ -kernel⁺-stable payoff distribution, the agents are sharing the valuation of the CS S .

compute- ϵ -kernel⁺(ϵ, u, S)

repeat

for each agent $i \in N$ do

$$\sigma_i \leftarrow \max_{\mathcal{C} \in \mathcal{C}_i} v(\mathcal{C}) - \sum_{k \in \mathcal{C}} u_k$$

$$\delta \leftarrow \max_{(i,j) \in N^2, u_i > v(\{i\})} \sigma_i - \sigma_j;$$

$$(i^*, j^*) \leftarrow \operatorname{argmax}_{(i,j) \in N^2, i \neq j, u_i > v(\{i\})} \sigma_i - \sigma_j;$$

if $(u_{j^*} - v(\{j^*\})) < \frac{\delta}{2}$ **then**

$$d \leftarrow u_{j^*} - v(\{j^*\})$$

else

$$d \leftarrow \frac{\delta}{2}$$

// d ensures individually rational payoffs

$$u_{i^*} \leftarrow u_{i^*} + d;$$

$$u_{j^*} \leftarrow u_{j^*} - d;$$

until $\frac{\delta}{v(S)} \leq \epsilon$;

Theorem 1 (Convergence). *Algorithm 2 converges to a payoff distribution in the kernel⁺.*

Proof. At each side payment, the strength of the agent that previously had the largest strength strictly decreases (property 3.1). Af-

ter the side payment, the value of the largest strength does not increase (property 3.2). By repeating payment from an agent with a low strength to the agent with the highest strength, the largest strength decreases. The process produces a sequence of monotonically decreasing maximum strengths. Since it is bounded below by zero the bound is reached in the limit. When the difference in strength between the agents is about to become zero, the side payments stop. Hence, the algorithm converges to a payoff distribution where all agents have the same strength (or cannot pay any agent as their payoff is equal to the value obtained when they form a singleton coalition) and hence, the corresponding payoff distribution is in the kernel⁺. \square

This theorem proves the existence of one payoff distribution to be in kernel⁺. Hence, it is guaranteed that if the agents adopt kernel⁺ as stability criterion, they will find an agreement, which was our initial goal. In addition, the fact that any two agents are in equilibrium provides a level of fairness.

3.4 Extension for games with externalities

It is also interesting to consider games with externalities, i.e., cooperative games where the value of a coalition depends on the CS. Sandholm and Lesser attribute these externalities to the presence of shared resources (if a coalition uses some resources, they will not be available to other coalitions), or when there are conflicting goals: non-members can move the world farther from a coalition's goal state [10]. Ray and Vohra in [9] state that a "recipe for generating characteristic functions is a minimax argument": the value of a coalition \mathcal{C} is the value \mathcal{C} gets when the non-members respond optimally so as to minimize the payoff of \mathcal{C} . This formulation acknowledges that the presence of other coalitions in the population may affect the payoff of the coalition \mathcal{C} . One example is a bargaining situation where agents need to negotiate over the same issues: when agents form a coalition, they can have a better bargaining position, as they have more leverage, and because the other party needs to convince all the members of the coalition. If the other parties also form a coalition, the bargaining power of the first coalition may decrease. Recently the topic has raised interest in AI. Rahwan et al. in [7] consider the problem of CS generation. Michalak et al. [5] tackle the problem of representing such games (they use a more compact description, still allowing efficient computation). Elkind et al. [2] consider CSs in weighted voting games.

To compute the excess of a coalition, one can take the maximum excess of that coalition in each CS containing it, and we can modify the definition of the kernel to take the externalities into account. The computation of such a kernel-stable payoff distribution would then require much more resources since there are a lot more coalition values to check. Since all the agents in the population are concerned with such a computation, the agents should cooperate to compute in a distributed fashion such a payoff distribution. Because of lack of space, we do not provide the details here. In the following we show an example showing how we could consider the idea of the kernel⁺ to compute a stable payoff distribution.

In Table 1, we provide an example of computation of a payoff distribution in the kernel⁺ for a game with externalities. We start with a payoff distribution where the valuation of the optimal CS (CS highlighted in Table 1) is shared equally between all agents. The strength of agent 2 is zero as there is no coalition where agent 2 is present which would have a positive marginal payoff. For agents 0, 1 and 3, the coalition they form in the optimal CS generates a marginal payoff of $3.392 - 3 \cdot 0.848 = 0.041$, and that is their strength. Hence,

agent 2 must make a side payment to either agent 0, 1, or 3. In the example, the payment is made to agent 3. As a result, the strength of agent 3 decreases, and the strength of agent 2 increases, but agent 2 still has the lowest strength. Agent 1 now has the largest strength and hence a side payment from agent 2 to agent 1 occurs next. The process iterates until the difference in strength is within $\frac{\epsilon}{3.392}$. Note that in the final outcome, agent 2 receives a payoff of 0.785, which is different from its payoff of 0.766 in the Kernel distribution, as 2 is forming a singleton coalition that generates 0.766. In this particular example, some of the valuation of the coalition $\{0, 1, 3\}$ is shared with another agent.

(a) Valuation Function

CS	value	coalition value
$\{\{0\}\{1\}\{2\}\{3\}\}$	2.114	$\{0\}$ 0.272 $\{1\}$ 0.123 $\{2\}$ 0.805 $\{3\}$ 0.915
$\{\{03\}\{1\}\{2\}\}$	1.403	$\{03\}$ 1.147 $\{1\}$ 0.041 $\{2\}$ 0.215
$\{\{0\}\{13\}\{2\}\}$	1.054	$\{0\}$ 0.363 $\{13\}$ 0.667 $\{2\}$ 0.023
$\{\{0\}\{1\}\{23\}\}$	2.503	$\{0\}$ 0.108 $\{1\}$ 0.874 $\{23\}$ 1.521
$\{\{02\}\{1\}\{3\}\}$	1.445	$\{02\}$ 0.141 $\{1\}$ 0.857 $\{3\}$ 0.448
$\{\{023\}\{1\}\}$	0.957	$\{023\}$ 0.089 $\{1\}$ 0.869
$\{\{02\}\{13\}\}$	1.730	$\{02\}$ 1.087 $\{13\}$ 0.642
$\{\{0\}\{12\}\{3\}\}$	2.018	$\{0\}$ 0.144 $\{12\}$ 0.984 $\{3\}$ 0.890
$\{\{03\}\{12\}\}$	1.923	$\{03\}$ 1.590 $\{12\}$ 0.333
$\{\{0\}\{123\}\}$	1.363	$\{0\}$ 0.769 $\{123\}$ 0.594
$\{\{01\}\{2\}\{3\}\}$	0.646	$\{01\}$ 0.142 $\{2\}$ 0.019 $\{3\}$ 0.485
$\{\{013\}\{2\}\}$	3.392	$\{013\}$ 2.626 $\{2\}$ 0.766
$\{\{01\}\{23\}\}$	1.256	$\{01\}$ 0.326 $\{23\}$ 0.930
$\{\{012\}\{3\}\}$	1.678	$\{012\}$ 1.623 $\{3\}$ 0.055
$\{\{0123\}\}$	1.786	$\{0123\}$ 1.786

(b) History of side payments

agent	0				1				2				3				payment	from	to
	Payoff				excess														
0	0.848	0.848	0.848	0.848	0.082	0.082	0.000	0.082	0.041	2	3								
1	0.848	0.848	0.807	0.889	0.041	0.041	0.000	0.041	0.021	2	3								
2	0.848	0.848	0.786	0.910	0.021	0.027	0.019	0.021	0.004	2	1								
3	0.848	0.852	0.782	0.910	0.017	0.023	0.023	0.017	0.003	0	2								
4	0.845	0.852	0.785	0.910	0.020	0.023	0.020	0.020	0.001	0	1								
5	0.844	0.853	0.785	0.910	0.020	0.021	0.020	0.020											

Table 1. Example side payments to reach a payoff distribution in the EK.

4 Conclusion

The traditional assumption in game theory is that the members of a coalition share the value of that coalition. In multiagent systems, it may be desirable to form a CS that maximizes social welfare. However, in the work that considers forming such optimal CS, the stability of the CS is not considered. One could use the nucleolus or the kernel for that optimal CS. However, this may not be fair for some agents.

In this paper, we discuss a solution that involves side-payment across coalitions. This is not the traditional assumption. We believe it could be an effective approach to add some fairness into the payoff distribution, and hence, improve the stability of the optimal CS. We use the maximum excess of a coalition to represent the strength of an agent, and the payoff is stable when all agents have the same strength. We proved that the new stability concept is non-empty since an algorithm always converge towards one stable payoff vector.

In future work, we would like to explore additional properties of the kernel⁺. For example, it would be interesting to study classes

of games where the kernel⁺ and the kernel differ. It would also be interesting to determine the computational complexity of finding a kernel⁺-stable payoff distribution for some specific representation. We are currently working on approximation schemes where the agents analyze only a subset of all CSs. Given the set of CSs analyzed, the agents are in equilibrium, which makes our approximation useful from a practical viewpoint.

In addition, we notice that there is a potential to further promote fairness by modifying Algorithm 2. As of now, the side-payment are bilateral and when multiple agents are candidates for a side payment, one is randomly chosen. In the example, for the first two side payments, agent 3 receives a side payments, but agents 0 and 1 could have received it as well, but in the end, agent 3 has a larger payoff than agents 0 and 1, which may be considered unfair. We are working to modify the payment scheme by considering payments between more than one agents. For example, in the current example, we would like agents 0, 1, and 3 to receive an equal payment from agent 2 in the first time step. Such modified payment schemes may restore parity to final payoffs to these agents and hence further enhance the fairness of our proposed scheme. As the convergence proof would remain unchanged, we are guaranteed that a fairer payoff distribution exists. We are working on finding a formal characterization of such payoff distributions.

REFERENCES

- [1] Robert J. Aumann and Jacques H Drèze, 'Cooperative games with coalition structures', *International Journal of Game Theory*, **3**(4), 217–237, (1974).
- [2] Edith Elkind, Georgios Chalkiadakis, and Nicholas R. Jennings, 'Coalition structures in weighted voting games', in *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08)*, pp. 393–397, (2008).
- [3] Donald B. Gillies, *Some theorems on n-person games*, Ph.D. dissertation, Department of Mathematics, Princeton University, Princeton, N.J., 1953.
- [4] Matthias Klusch and Onn Shehory, 'A polynomial kernel-oriented coalition algorithm for rational information agents', in *Proceedings of the Second International Conference on Multi-Agent Systems*, pp. 157 – 164. AAAI Press, (December 1996).
- [5] Tomasz Michalak, Talal Rahwan, Jacek Sroka, Andrew Dowell, Michael Wooldridge, Peter McBurney, and Nicholas R. Jennings, 'On representing coalitional games with externalities', in *Proceedings of the 10th ACM conference on Electronic Commerce 09 (EC'09)*, (2009).
- [6] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, The MIT Press, 1994.
- [7] Talal Rahwan, Tomasz Michalak, Nicholas R. Jennings, Michael Wooldridge, and Peter McBurney, 'Coalition structure generation in multi-agent systems with positive and negative externalities', in *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI-09)*, (2009).
- [8] Talal Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, and Andrea Giovannucci, 'An anytime algorithm for optimal coalition structure generation', *Journal of Artificial Intelligence Research*, **34**, 521–567, (2009).
- [9] Debraj Ray and Rajiv Vohra, 'A theory of endogenous coalition structures', *Games and Economic Behavior*, **26**, 286–336, (1999).
- [10] Tuomas Sandholm and Victor R. Lesser, 'Coalitions among computationally bounded agents', *AI Journal*, **94**(1–2), 99–137, (1997).
- [11] Onn Shehory and Sarit Kraus, 'Feasible formation of coalitions among autonomous agents in nonsuperadditive environments', *Computational Intelligence*, **15**, 218–251, (1999).
- [12] Richard Edwin Stearns, 'Convergent transfer schemes for n-person games', *Transactions of the American Mathematical Society*, **134**(3), 449–459, (December 1968).