

The Good, The Bad and The Cautious: Safety Level Cooperative Games

Yoram Bachrach¹ Maria Polukarov² Nicholas R. Jennings²

¹Microsoft Cambridge, United Kingdom

yobach@microsoft.com

²University of Southampton, United Kingdom

{mp3,nrj}@ecs.soton.ac.uk

Abstract

We study *safety level coalitions* in competitive games. Given a normal form game, we define a corresponding *cooperative* game with transferable utility, where the value of each coalition is determined by the safety level payoff it derives in the original—*non-cooperative*—game. We thus capture several key features of agents’ behavior: (i) the possible monetary transfer among the coalition members; (ii) the solidarity of the outsiders against the collaborators; (iii) the need for the coalition to optimize its actions against the worst possible behavior of those outside the coalition. We examine the concept of safety level cooperation in *congestion games*, and focus on computing the value of coalitions, the core and the Shapley value in the resulting safety level cooperative games. We provide tractable algorithms for *anonymous* cooperative games and for safety level cooperative games that correspond to *symmetric* congestion games with singleton strategies. However, we show hardness of several problems such as computing values in games with multi-resource strategies or asymmetric strategy spaces.

1 Introduction

Game theory analyzes the interactions of multiple self-motivated, rational agents, each wishing to maximize its own profits from the interaction. In such situations, an agent may not follow any centrally “prescribed” behavior if deviating from it improves its utility, and so the study of *stable* outcomes in games is the central topic in game theory. In *non-cooperative* (or, *competitive*) games, where agents take individual actions, the prominent stability concept is *Nash equilibrium*—a strategy profile where no agent has a beneficial unilateral deviation from. However, it does not take into account collective deviations by groups of agents; to this end, the notion of *strong equilibrium* [2]—a strategy profile with no profitable deviations by subsets of agents—extends Nash equilibrium to capture collective strategic behavior. *Cooperative* (or, *coalitional*) games consider how coalitions of agents can cooperate, focusing on how the utility is distributed among the agents. The *core* [7] of a coalitional game is composed by all possible ways the agents can share the utility so that no subset of the agents has an incentive to deviate.

In a non-cooperative game, agents act not only selfishly but also independently, and base their decisions on their individual interests. Thus, even when they are able to coordinate their actions, each agent has to enjoy the outcome of cooperation, individually. However, in many scenarios, traditionally modeled as non-cooperative games (e.g., auctions, network formation games, routing and congestion games), a group (coalition) of agents may well need to *jointly* decide on the collective action and, also, make monetary transfers to *share* the gains the coalition obtains. This requires applying tools from cooperative game theory to non-cooperative domains (which rarely overlap with typical cooperative scenarios).

To this end, Hayrapetyan et al [8] proposed a framework for modelling coalitions in the context of *congestion games*. In a classic congestion game [13], a set of agents shares a set of resources, and

an agent’s strategy is to choose a subset of resources to use, so that the sum of congestion-dependent costs over its selected resources is minimized. In [8], the authors assume that the agents may collude, aiming to maximize their collective welfare. Their model also allows monetary transfers. However, they address cooperation from a somewhat different perspective than cooperative game theory, focusing on the possible negative effect of collusion on the social welfare. Another work [3] deals with formation of coalitions in VCG multi-unit auctions. However, there, it is assumed that the outsiders of a coalition would always truthfully reveal their individual preferences and play according to them.

In contrast, this paper studies the stability and computational aspects of *safety level* coalitions in competitive games. As opposed to a standard cooperative game, the utility of a coalition in such settings depends not only on the action the members take, but also on the actions taken by the non-members. In the worst case, the outsiders may decide to punish the members of the coalition and take actions that minimize the collective utility of the collaborators. A coalition then may decide to maximize its (total) utility under the worst case action of the non-members—we call this a *joint safety level* strategy. In other words, in scenarios we consider the collaborators are “good” to each other by that they cooperatively choose actions and share gains, and the non-collaborators may behave as “bad” adversaries who reduce the collaborators’ utility. If the collaborators take this view and consider non-collaborators as adversaries, they must be “cautious” and prepare for the worst-case choice of the non-collaborators, by using their joint safety level strategy. To achieve this cooperation, the agents can agree on monetary transfers in the form of an enforceable contract, allowing them to distribute the gains in a certain way. Solution concepts such as the core and the Shapley value [17] can be used to predict what transfers would occur.

We apply the concept of safety level cooperation to the class of congestion games. These games, in which selfish agents strategically choose from a common set of resources and derive individual utilities that depend on the total congestion on each resource, are fundamental to a wide range of applications. Examples include resource and task allocation, firm competition for production processes, routing problems, network design, and other kinds of resource sharing scenarios in distributed systems [12, 13, 14]. Such games are known to have Nash equilibria in pure strategies [13], and some of their restricted classes are known to have strong equilibria in pure strategies [9, 15]. Also, much recent effort in algorithmic game theory has gone into study of interesting subclasses of congestion games that are computationally tractable [1, 10]. One important subclass is *resource selection games*, in which each agent is restricted to choosing a single resource. The resource selection games are in a particular focus of our work.

Our key results distinguish between *symmetric* settings where agents choose strategies from a common space, and *asymmetric* ones where each agent has its own collection of strategies. While in the non-cooperative context, both symmetric and asymmetric models are anonymous, asymmetric models lose anonymity when monetary transfers are allowed. For anonymous settings, we show that testing the core (non-)emptiness, constructing a core imputation and testing whether an imputation is in the core are in P when the computation of the coalitional values in the game is in P ; we also show that the Shapley value is in the core if it is not empty, and can be computed in polynomial time. These results hold for *all* anonymous cooperative games¹—not only those based on safety level coalitions. For congestion games, we show that computing a coalition’s value is in P for singleton strategies and is NP-hard for multiple-resource strategies, while for non-anonymous settings computing the value of even a singleton or the grand coalition are NP-hard.

2 Preliminaries

We start with necessary game-theoretic definitions, and define congestion games.

A *non-cooperative game in normal form* is given by an agent set $\mathbf{N} = \{1, \dots, N\}$, and for each agent $i \in \mathbf{N}$, a strategy space \mathbf{S}_i of its pure strategies and a payoff function $U_i : \times_{i \in \mathbf{N}} \mathbf{S}_i \rightarrow \mathbb{R}$ specifying the

¹Not to be confused with “coalitional games in open anonymous environments” [21], which proposes *anonymity-proof* solutions (with respect to the core and the Shapley value) which are robust under “false name” manipulations by the agents. In contrast, we refer to cooperative games where the characteristic function is not “sensitive” to the agents’ identities—that is, all coalitions of the same size get equal values.

award an agent gets at the outcome of the game. Denote by \mathbf{S}_C the set of partial strategy profiles of a subset of agents $C \subseteq \mathbf{N}$, and by $\mathbf{S}_{-C} = \mathbf{S}_{\mathbf{N} \setminus C}$ the set of strategy combinations of all the agents outside C ; for a single agent $i \in \mathbf{N}$, denote $\mathbf{S}_{-i} = \mathbf{S}_{\mathbf{N} \setminus \{i\}}$. A strategy profile $s \in \mathbf{S}$ is a *Nash equilibrium* if for each agent $i \in \mathbf{N}$ and for each its strategy $s'_i \in \mathbf{S}_i$ the following holds: $U_i(s) \geq U_i(s_{-i}, s'_i)$. A strategy profile is a *strong Nash equilibrium* if it is stable against deviations by coalitions: for any $C \subseteq \mathbf{N}$ and $s'_C \in \mathbf{S}_C$, there exists $i \in C$ such that $U_i(s) \geq U_i(s_{-i}, s'_i)$. The *safety level strategy* for agent $i \in \mathbf{N}$, s_i^{SL} , is the strategy maximizing its *guaranteed* utility, no matter what the other agents play: $s_i^{SL} \in \arg \max_{s_i \in \mathbf{S}_i} \min_{s_{-i} \in \mathbf{S}_{-i}} U_i(s_i, s_{-i})$.

It is crucial that an agent's payoff is determined by the profile of strategies played by all participants. Some utility functions ignore the identities of the agents, and only take into account *how many times* each strategy is played. These settings, where identities are irrelevant, are *anonymous*². Given a set of strategies $\mathbf{S} = \{1, \dots, S\}$, a strategy $s \in \mathbf{S}$ and an agent $i \in \mathbf{N}$, the utility of i playing s in an anonymous game is a function mapping the set of partitions $\{(x_1, \dots, x_S) \mid x_j \in \{1, \dots, N\}, \sum_{j=1}^S x_j = N - 1\}$ to real numbers. A related important subclass is that of *symmetric games*, where the payoffs for playing a particular strategy are the same for different agents and depend only on the other strategies employed: if one can change the identities of the agents without changing the payoffs to the strategies, then a game is symmetric. Formally, a game with strategy spaces given by $\mathbf{S}_1 = \dots = \mathbf{S}_N = \mathbf{S}$ is symmetric if for any permutation π over \mathbf{N} and an agent $i \in \mathbf{N}$, we have $U_i(s_1, \dots, s_i, \dots, s_N) = U_{\pi(i)}(s_{\pi(1)}, \dots, s_{\pi(i)}, \dots, s_{\pi(N)})$, where $s_j = s_{\pi(j)}$ for $j = 1, \dots, N$.

Our analysis is also based on cooperative game theory, so we briefly review relevant notions. A *transferable utility coalitional* (or, *cooperative*) *game* is composed of a set \mathbf{N} of N agents, and a characteristic function $v : 2^{\mathbf{N}} \rightarrow \mathbb{R}$ mapping any subset (coalition) of agents to a real value, indicating the total utility these agents achieve together. We denote the set of all the agents except i as $\mathbf{N}_{-i} = \mathbf{N} \setminus \{i\}$. A coalitional game is *monotone* if for all coalitions $C' \subseteq C$ it holds that $v(C') \leq v(C)$, and is *super-additive* if for all *disjoint* coalitions $A, B \subseteq \mathbf{N}$ we have $v(A) + v(B) \leq v(A \cup B)$. In super-additive games, it is always worthwhile for two sub-coalitions to merge, so eventually the *grand coalition* of all the agents will form.

The characteristic function only indicates the total gains a coalition can achieve, but does not specify how these gains are distributed among the agents who formed it. An *imputation* (p_1, \dots, p_N) defines a division of the gains of the grand coalition among its agents, where $p_i \in \mathbb{R}$, such that $\sum_{i=1}^N p_i = v(\mathbf{N})$. We call p_i the payoff of agent i , and denote the payoff of a coalition C as $p(C) = \sum_{i \in C} p_i$. A basic requirement for a good imputation is *individual rationality*: for any agent $i \in \mathbf{N}$, $p_i \geq v(\{i\})$ (otherwise, this agent is incentivized to work alone). Similarly, we say a coalition B *blocks* imputation (p_1, \dots, p_N) if $p(B) < v(B)$. If a blocked payoff vector is chosen, the coalition is somewhat unstable. The most prominent solution concept based on such *stability* is the core [7]. The *core* is the set of all imputations (p_1, \dots, p_N) not blocked by any coalition, so that for any coalition $C \subseteq \mathbf{N}$ holds $p(C) \geq v(C)$.

Another solution concept is the Shapley value [17] which focuses on *fairness*, rather than stability. The Shapley value fulfills several important fairness axioms [17, 20] and has been used to fairly share gains or costs. For each agent, it depends on the agent's marginal contribution to possible coalition permutations. We denote by π a permutation (ordering) of the agents, and by Π the set of all possible such permutations. Given a permutation $\pi = (i_1, \dots, i_N) \in \Pi$, the *marginal worth vector*, $m^\pi(v) \in \mathbb{R}^N$, is defined by $m_{i_1}^\pi = v(\{i_1\})$ and $m_{i_k}^\pi(v) = v(\{i_1, i_2, \dots, i_k\}) - v(\{i_1, i_2, \dots, i_{k-1}\})$ for $k > 1$. The convex hull of all the marginal vectors is called the *Weber Set*. It has been shown [19] that the Weber set of any game contains its core. The *Shapley value* is the centroid of the marginal vectors: $\phi(v) = \frac{1}{N!} \sum_{\pi \in \Pi} m^\pi(v)$.

We analyze the core and the Shapley value of cooperative games that arise when considering safety-level coalitions in a given non-cooperative setting, and demonstrate this approach on congestion games. In a *congestion game* (CG) [13], every agent has to choose from a finite set of resources. The utility of an agent from using a particular resource depends on the number of agents using it, and its total utility is the sum of utilities on its used resources. Formally, a congestion game $\Gamma = (\mathbf{N}, \mathbf{R}, (u_r(\cdot))_{r \in \mathbf{R}})$ is described by the following components: a set $\mathbf{N} = \{1, \dots, N\}$ of agents; a set $\mathbf{R} = \{r_1, \dots, r_R\}$

²See [6] for results on approximating equilibria in anonymous games.

of resources; an assignment $u_r : \{1, \dots, N\} \rightarrow \mathbb{R}$, $r \in \mathbf{R}$, of resource utility functions, where for any resource $r \in \mathbf{R}$, $u_r(k)$ is the resource utility (cost) for r when the total number of users of r is k . Each agent i is allowed to choose a (non-empty) bundle of resources $B \subseteq 2^{\mathbf{R}}$, from a certain set $\mathbf{S}_i = \{B_1^i, \dots, B_{s_i}^i\}$ of allowed bundles (where each $B_j^i \subseteq \mathbf{R}$). We denote by $s_i \in \mathbf{S}_i$ the strategy (set of resources) chosen by agent i . Every N -tuple of strategies—a *strategy profile*— $s = (s_i)_{i \in \mathbf{N}}$ corresponds to an R -dimensional congestion vector $h(s) = (h_r(s))_{r \in \mathbf{R}}$ where $h_r(s)$ is the number of agents who select resource r (we simply write h_r when it's clear what profile we refer to). The utility of i from s is: $U_i(s) = \sum_{e \in s_i} u_r(h_r(s))$. A congestion game is a *resource selection game* (RSG) if the strategy space of every agent corresponds to a set of singletons. That is, agent i chooses a single resource from the given set, and its payoff from a strategy profile $s = (s_i)_{i \in \mathbf{N}}$ is given by $U_i(s) = u_{s_i}(h_{s_i}(s))$.

Remark 1. *In a congestion game an agent's utility only depends on the numbers of agents choosing each resource but not on their identities, so congestion games are anonymous. Since the utility from each particular resource is the same for each of its users, the utility any agent gets from playing a particular strategy depends only on the other strategies selected, but not on who has chosen them. Thus a congestion game is symmetric if (and only if) all agents in the game have identical strategy spaces. We refer to symmetric congestion and resource selection games as SCGs and SRSGs, respectively.*

Congestion games always have a pure strategy Nash equilibrium [13]. Resource selection games with monotone utility functions also admit strong equilibria [9]. In fact, in RSGs with decreasing utilities, any Nash equilibrium is strong. However, we show that coalitional stability is no longer guaranteed if utility transfers are allowed.

3 Safety Level Cooperative Games

Let $\Gamma = (\mathbf{N}, (\mathbf{S}_i)_{i \in \mathbf{N}}, (U_i)_{i \in \mathbf{N}})$ be a normal-form game, where \mathbf{N} is the set of agents, and \mathbf{S}_i and U_i denote, respectively, strategy spaces and utility functions of *individual* agents. In this context, we are interested in scenarios where it makes sense to the agents to form coalitions and coordinate their actions to optimize their *collective* gains. We take a *safety level* approach to analyzing gains of a coalition. Thus, we assume that the coalition members attempt to maximize the minimal utility they would get under any strategy choices of the non-members. We model coordination in the underlying normal-form game as a coalitional game, where coalitional values are determined by the safety-level payoffs of each coalition. To this end, we first extend the notion of a safety level to coalitional payoffs.

For coalition $C \subseteq \mathbf{N}$ and strategy profile $s = (s_i)_{i \in \mathbf{N}}$, let $U_C(s) = \sum_{i \in C} U_i(s)$ be the total utility C achieves under s . The utility of a coalition depends not only on the strategies chosen by its members, but also on the choices of the non-members. Let $B = \mathbf{N} \setminus C$ denote the set of non-members. Now, a profile s can be written as $s = (s_B, s_C)$, where $s_C = (s_i)_{i \in C}$ and $s_B = (s_j)_{j \in B}$ are partial strategy profiles. Given the non-members' strategy s_B , the coalition could optimize for the total value it can achieve, by choosing $s_C^* \in \arg \max_{s_C \in \mathbf{S}_C} U_C(s_B, s_C)$, where $\mathbf{S}_C = \times_{i \in C} \mathbf{S}_i$ is the set of coalitional strategies of C . This choice maximizes C 's utility for a *specific* strategy profile of B . What should the coalition C do without knowing how the non-members would behave? Staying on the “safe” side, C can optimize the utility guaranteed to it, no matters what the outsiders do, by maximizing its *safety level*, the worst case utility the coalition obtains under all possible actions of the non-members. The safety level of C when it chooses s_C is: $U_C^{SL}(s_C) = \min_{s_B \in \mathbf{S}_B} U_C(s_B, s_C)$, and the *safety level strategy* of a coalition C is the coalitional strategy $s_C^* \in \mathbf{S}_C$ that maximizes the safety level:

$$s_C^* \in \arg \max_{s_C \in \mathbf{S}_C} U_C^{SL}(s_C) = \arg \max_{s_C \in \mathbf{S}_C} \left(\min_{s_B \in \mathbf{S}_B} U_C(s_B, s_C) \right)$$

The *safety level value* of C is the minimal utility it would get when using its safety level strategy:

$$U_C^* = \min_{s_B \in \mathbf{S}_B} U_C(s_B, s_C^*) = \max_{s_C \in \mathbf{S}_C} \left(\min_{s_B \in \mathbf{S}_B} U_C(s_B, s_C) \right)$$

A coalition's safety level value is the utility the coalition can guarantee to itself *as a whole* when its members cooperate. A key challenge determining how the members would share this value. To answer this question, we define a *safety level cooperative game* (SLC-game) for Γ as follows:

Definition 1 (Safety Level Cooperative Game). *Given a (normal-form) game $\Gamma = (\mathbf{N}, (\mathbf{S}_i)_{i \in \mathbf{N}}, (U_i)_{i \in \mathbf{N}})$ with agent set \mathbf{N} , strategy space \mathbf{S}_i and utility function U_i for each $i \in \mathbf{N}$, the induced safety level cooperative game (SLC-game) is a cooperative game over the same set \mathbf{N} of agents, where the characteristic function is the safety level value of coalitions in Γ : for each $C \subseteq \mathbf{N}$, $v(C) = U_C^*$.*

We will write SLC^Γ to indicate that an SLC-game is induced by a game Γ . However, regardless of their underlying games Γ , all SLC-games have the following basic property:

Lemma 1. *The SLC-games are monotonically increasing.*

Proof. We need to show that for any C', C such that $C' \subseteq C$ we have $v(C') \leq v(C)$. Intuitively, as C includes more agents than C' and the agents in $D = C \setminus C'$ are coalition members for C and outsiders for C' , so they “help” the members of C and “punish” the members of C' . Hence, the safety level value of a larger coalition is greater than that of a smaller one. Formally, denote $B = \mathbf{N} \setminus C$, so $\mathbf{N} \setminus C' = B \cup D$. We have:

$$\begin{aligned} v(C') = U_{C'}^* &= \max_{s_{C'} \in \mathbf{S}_{C'}} \left(\min_{s_{B \cup D} \in \mathbf{S}_{B \cup D}} U_{C'}(s_{B \cup D}, s_{C'}) \right) \leq \max_{s_{C'} \in \mathbf{S}_{C'}} \left(\min_{s_B \in \mathbf{S}_B} U_{C'}(s_B, s_D, s_{C'}) \right) \\ &\leq \max_{s_{C'} \in \mathbf{S}_{C'}} \left(\min_{s_B \in \mathbf{S}_B} U_C(s_B, s_D, s_{C'}) \right) \leq \max_{s_C \in \mathbf{S}_C} \left(\min_{s_B \in \mathbf{S}_B} U_C(s_B, s_C) \right) = U_C^* = v(C) \end{aligned}$$

□

4 Safety Level Coalitions in Congestion Games

We now analyze safety level coalitions in congestion games and resource selection games. We also make a distinction between symmetric settings where agents derive strategies from a common space, and asymmetric settings where each agent has its own collection of strategies. While in the non-cooperative context, both symmetric and asymmetric models are anonymous, asymmetric models lose anonymity when monetary transfers are allowed. We show anonymous and non-anonymous SLC-games have different computational properties.

Consider a congestion game with agents \mathbf{N} and resources \mathbf{R} with resource utility functions $u_r(\cdot)$ for $r \in \mathbf{R}$, and a coalition $C \subseteq \mathbf{N}$. For any strategy profile $s = (s_i)_{i \in \mathbf{N}}$, the congestion on each resource is $h(s) = (h_r(s))_{r \in \mathbf{R}}$, and we can compute the utility $U_i(s)$ for each agent i . C 's total utility under s is $U_C(s) = \sum_{i \in C} U_i(s) = \sum_{i \in C} \sum_{r \in s_i} u_r(h_r(s))$. Denote the number of C 's members who use a resource r at a strategy profile s as $h_r^C(s) = |\{i \in C \mid r \in s_i\}|$. We can write: $U_C(s) = \sum_{r \in \mathbf{R}} h_r^C(s) \cdot u_r(h_r(s))$. The coalitional value of C in the corresponding SLC-game is:

$$v(C) = U_C^* = \max_{s_C \in \mathbf{S}_C} \left(\min_{s_B \in \mathbf{S}_B} \sum_{r \in \mathbf{R}} h_r^C(s) \cdot u_r(h_r(s)) \right)$$

Recall our notation of (S)CG and (S)RSG for (symmetric) congestion and resource selection games. Note that $SRSGs \subseteq RSGs \subseteq CGs$ and $SRSGs \subseteq SCGs \subseteq CGs$. Similar inclusions hold for the corresponding safety level cooperative games, induced by (S)CGs and (S)RSGs.

4.1 Anonymous cooperative games

We consider the properties of SLC-games induced by the subclass of symmetric congestion games, where all agents use a common set of strategies. We start with Lemma 2 below showing that these games satisfy what we call *anonymity*, naturally defined as follows: we say a cooperative game is *anonymous* if any two agents are equivalent—i.e., for every two agents $i \neq j$ and any coalition C such that $i \notin C$ and $j \notin C$ we have $v(C \cup \{i\}) = v(C \cup \{j\})$.

Lemma 2. All SLC^{SCG} -games are anonymous.

Proof. Consider a coalition C that contains neither i nor j . Since both i and j have identical strategy spaces, we get the same sets for min and max operators when computing coalitional safety level values of $C \cup \{i\}$ and $C \cup \{j\}$. \square

Now, for *all* anonymous cooperative games (and for SLC^{SCG} -games in particular) we have that the Shapley value vector (i) can be found in polynomial time and (ii) belongs to the core if the core of a given game is not empty (see the following Lemmas 3 and 4; for proofs, see Appendix A):

Lemma 3 (Core of Anonymous Games). *Let v be an anonymous cooperative game over N agents \mathbf{N} , with a non-empty core. Denote $q = \frac{v(\mathbf{N})}{N}$. Then the symmetric payoff distribution (q, q, \dots, q) is an imputation in the core.*

Lemma 4 (Shapley Value of Anonymous Games). *Let v be an anonymous cooperative game over N agents \mathbf{N} . Denote $q = \frac{v(\mathbf{N})}{N}$. Then the Shapley value is the symmetric payoff distribution (q, q, \dots, q) . If the core exists, then the Shapley value is in the core.*

Lemmas 3 and 4 require the non-emptiness of the core. However, some safety level cooperative games have non-empty cores and others have empty cores (see Examples 1 and 2 below). Such examples can occur even among the restricted class of SLC-games induced by symmetric, monotone resource selection games, which always possess strong equilibria, highlighting the difference between the cooperative safety level cooperative game's core and strong equilibrium.

We now introduce additional notation. In an anonymous game all agents are equivalent so the value of a coalition only depends on the *number* of agents in the coalition and not their identities. Thus, we can write the characteristic function v as a function mapping the size of a coalition to its value, so $v : \{0, 1, \dots, N\} \rightarrow \mathbb{R}$. We use the standard convention that $v(0) = 0$.

Example 1 (there exist SLC^{CG} -games with non-empty core). *The example is a SLC^{SRSG} -game with N agents and R resources with identical, constant resource utility functions $u_r(k) = x \in \mathbb{R}$ for any $r \in \mathbf{R}$, $k = 1, \dots, N$. We show that any such game has a non-empty core.*

Proof. Note that the value of any coalition in this domain, no matter what the non-members do, only depends on the size of the coalition, so $v(k) = xk$. Thus, the simple payoff vector $p = (x, \dots, x)$ is in the core, since given any coalition C of size $|C|$ we have $p(C) = x|C| = v(C)$, and all the core conditions hold. \square

Example 2 (SLC^{SRSG} -games may have empty core). *Consider an SLC^{SRSG} -game with $N = 3$ agents and two resources $\{a, b\}$ with identical resource utility functions $u_r(1) = 2; u_r(2) = u_r(3) = 1$ for $r = a, b$. The core of this game is empty.*

Proof. We have $v(0) = 0$. Now compute $v(1)$, the safety level of a *single* agent (out of 3 agents in total). No matter which resource, a or b , the agent chooses, the worst case outcome is when the other 2 agents also choose the same resource, giving the agent a utility of $u_r(3) = 1$; thus, we have $v(1) = 1$. Now consider the safety level of 2 agents (out of 3). They can either choose to have both of them using the same resource, or to use a different resource each. If they both are on the same resource, the worst case action of the remaining agent is to also join that resource, and the utility of the coalition is $2u_r(3) = 2$. If the collaborators choose different resources, any choice of the remaining agent results in having 2 agents (one member and one non-member of the coalition) on one resource and a single coalition member on the other resource, resulting in a utility of $u_r(2) + u_r(1) = 2 + 1 = 3$ for the coalition. Thus, the safety level of any pair of agents is $v(2) = 3$. A coalition of 3 agents is the grand coalition, whose best choice is to assign 2 agents on one resource, and 1 agent on the other resource, and so $v(3) = u_r(1) + 2u_r(2) = 2 + 2 = 4$. Thus, the characteristic function of this SLC-game is given by $v(0) = 0, v(1) = 1, v(2) = 3, v(3) = 4$. Due to Lemma 3, if the game has a non-empty core, the imputation $p = (\frac{4}{3}, \frac{4}{3}, \frac{4}{3})$ should be in the core. However, under this imputation the payoff for any two agents is less than the value of a coalition of the

pair: i.e., $p(\{1, 2\}) = \frac{8}{3}$, but $v(\{1, 2\}) = v(2) = 3 > \frac{8}{3} = p(\{1, 2\})$, which violates the core constraints. Hence, the core is empty. \square

Moreover, restricting or expanding the sets of the agents' allowed strategies may cause the core to change from being empty to being non-empty and vice versa:

Example 3 (Strategy Sets and the Core). *Consider the game with 3 agents and 2 resources $\{a, b\}$ from the previous example, where the resource utility function is given by $u_r(1) = 2; u_r(2) = u_r(3) = 1$ for $r = a, b$. The core of this game is empty. Now add a third resource c with a constant utility of $u_c(k) = 10$ for $k = 1, 2, 3$, and expand each agent's strategy set to allow selecting $\{c\}$. The resulting game is anonymous, with characteristic function $v(1) = 10, v(2) = 20, v(3) = 30$, and its core is not empty: the imputation $p = (10, 10, 10)$ is in the core. On the other hand, if we take this new game, and restrict each agent's strategy set to allow selecting only $\{a\}$ or $\{b\}$, we obtain the original game with an empty core. Thus, extending strategy sets makes the core non-empty, and restricting them may empty it.*

Now, consider the game with 3 resources $\{a, b, c\}$, where again $u_r(1) = 2; u_r(2) = u_r(3) = 1$ for $r = a, b$, but $u_c(k) = 0.1$ for $k = 1, 2, 3$. If the agents are restricted to choosing only c , i.e. $\mathbf{S}_1 = \mathbf{S}_2 = \mathbf{S}_3 = \{\{c\}\}$, we have an anonymous game where $v(1) = 0.1, v(2) = 0.2, v(3) = 0.3$ which has a non-empty core as $p = (0.1, 0.1, 0.1)$ belongs to it. If we extend the strategy sets to also include a and b , so that $\mathbf{S}_1 = \mathbf{S}_2 = \mathbf{S}_3 = \{\{a\}, \{b\}, \{c\}\}$, we get the game where $v(1) = 1, v(2) = 3, v(3) = 3$, whose core is empty. Thus, extending strategy sets may make the core empty, and shrinking them makes it non-empty.

Remark 2. *Based on the above examples, one can see that the non-cooperative and cooperative concepts of coalitional stability are rather different. While strong Nash equilibria always exist for (monotone) resource selection games, the core of their corresponding SLC-games may be empty. The reason for that is the following: while for any coalition there could be no deviation guaranteeing a better payoff to any of the deviators, there might exist a coalition that can improve its total welfare—that is, even if some agents may obtain worse individual utilities after the deviation, this loss will be covered by the gains their co-deviators get.*

In light of the above observations, testing the (non-)emptiness of the core in safety level cooperative games is an important issue. It follows from the next Theorem 1 regarding anonymous cooperative games, that for SLC^{SCG} -games this can be done efficiently if the computation of coalitional values is easy; moreover, in this case, the construction of a core imputation and verification if a given imputation is in the core are also computationally efficient:

Theorem 1 (Core Computation in Anonymous Games). *In anonymous cooperative games, if computing the value of any coalition can be performed in polynomial time, then the following problems are in P : testing for core emptiness, constructing a core imputation (if one exists) and testing if an imputation p is in the core.*

Proof. In anonymous games the characteristic function is given as $v : \{0, 1, \dots, N\} \rightarrow \mathbb{R}$ —the function that maps the size of a coalition to its value. This representation is simply a table, containing N numbers: therefore, if computing the value of each coalition can be performed in polynomial time, then finding the characteristic function is also so.

To fulfil the core constraints, the following must hold for an imputation p : $\sum_{i=1}^N p_i = v(\mathbf{N}) = v(N)$, and $\forall C, p(C) \geq v(C)$. Consider testing whether an imputation p satisfies this. It is easy to check if $\sum_{i=1}^N p_i = v(\mathbf{N}) = v(N)$. However, testing the condition $\forall C, p(C) \geq v(C)$ seemingly requires 2^N similar tests. Order the agents according to their payment, so that $p_{i_1} \leq p_{i_2} \leq \dots \leq p_{i_N}$. Denote by C_k the coalition $C = \{i_1, i_2, \dots, i_k\}$. Note that if the core constraint $p(C) \geq v(C)$ holds for $C_k = \{i_1, i_2, \dots, i_k\}$, it must also hold for any coalition of size k , as C_k is the minimally paid coalition of size k . Thus, to test if p is a core imputation, it is enough to test whether $p(C_k) \geq v(C_k)$ for $k \in \{1, 2, \dots, N\}$. If the core constraints hold for all C_1, \dots, C_N , they hold for any coalition C , and

if they do not, we have a violated constraint. Since there are only N such checks, this can be done in polynomial time.

Now consider testing for core-emptiness and constructing a core imputation. Due to Lemma 3, if the core is non-empty, the symmetric imputation (q, q, \dots, q) where $q = \frac{v(\mathbf{N})}{N}$ must be in the core. Since q can be computed in polynomial time, this imputation can also be computed in polynomial time. We can then test whether it is in the core. If it is in the core, we have constructed a core imputation, and if it is not in the core, then the core is empty. \square

However, the computation of safety levels and the values of coalitions in SLC-games, can be a non-trivial task, as the safety level strategies of the agents are not even robust to small changes of game parameters. For instance, we show that even changing only the total number of agents can result in very different safety level strategies, even in simple anonymous settings.

Example 4 (Number of Agents and Safety Level Strategies). *Consider an SLC^{SRSG}-game with two resources $\{a, b\}$ with resource utility functions given by $u_a(k) = \epsilon$ for $k = 1, \dots, 5$ an a small positive ϵ , and $u_b(k) = 1 \forall k = 1, \dots, 4$, $u_b(5) = 0$. Assume there are $N = 4$ agents playing the game and compute the value of a coalition C of 3 agents (out of 4). Since $u_b(\cdot)$ is constant up to congestion of 4, any agent in C who chooses b is guaranteed a utility of 1 on that resource. On the other hand, any agent in C who chooses a only gets a utility of ϵ on that resource. Thus $v(3) = 3 \cdot 1 = 3$, and the safety level strategy of C is to have all its agents choosing the resource b . Now, consider the same resources and resource utility functions when there are $N = 5$ agents, and consider again a coalition C of 3 agents. If C places all agents in b , a possible strategy for the remaining 2 agents is to both join b , resulting in a total utility of $3u_b(5) = 3 \cdot 0 = 0$ for the coalition. Alternatively, the coalition can have 2 agents using a and 1 agent using b . For this strategy in \mathbf{S}_C , any strategy in $\mathbf{S}_{\mathbf{N} \setminus C}$ of the remaining 2 agents results in all agents on a getting a utility of ϵ and all those on b getting 1, resulting in a total coalitional utility of $1 \cdot \epsilon + 2 \cdot 1 = 2 + \epsilon$. This is the safety level strategy for the coalition C , so $v(3) = 2 + \epsilon$.*

Nevertheless, for SLC^{SRSG}-games we can compute a coalition's value in polynomial time:

Theorem 2. *For SLC^{SRSG}-games, computing safety level strategies and values is in P.*

Proof. We provide a dynamic programming algorithm. Given an SLC^{SRSG}-game with R resources, for any $k = 1, \dots, R$ let v_k denote the k -subgame, played on the first k resources: that is, v_k is the restriction of the original game where the agents are only allowed to select one of the first k resources—i.e., for each $i \in \mathbf{N}$ we have $\mathbf{S}_i = \{r_1, r_2, \dots, r_k\} \subseteq \mathbf{R}$. Note that v_k is also an SLC^{SRSG}-game. We denote by $v_{i,j,k}$ the value of a coalition of i agents in the k -subgame with $i + j$ agents: to compute $v_{i,j,k}$ we must find a safety level strategy for a coalition of i agents when there are additional j non-members, and the agents are only allowed to select one of the first k resources. We prove that the following recursive formula holds:

$$v_{i,j,k} = \max_{p \in \{1, 2, \dots, i\}} \left(\min_{q \in \{1, 2, \dots, j\}} (v_{i-p, j-q, k-1} + p \cdot u_k(p+q)) \right)$$

Consider a coalition C of i agents, who are currently using the safety level strategy in the $(k-1)$ -subgame with additional j agents. The coalition assigns c_x agents to use resource x (where $x \leq k-1$), so that $\sum_{x=1}^{k-1} c_x = i$. The worst case response of the non-members in this subgame is assigning b_x resources to use resource x , (where $x \leq k-1$), so that $\sum_{x=1}^{k-1} b_x = j$. We can describe a strategy for C in the k -subgame in terms of *moving* some p agents from the first $k-1$ resources and assigning them to the resource k . Any strategy for C in the k -subgame can be described as having $p \leq i$ coalition members using resource k , and a partition of the $i-p$ remaining agents to the first $k-1$ resources (which is a strategy for a coalition C' of $i-p$ agents in the $(k-1)$ -subgame), for some choice of $p \leq i$. Each such a partial strategy profile s_C implies a response from the non-members $\mathbf{N} \setminus C$ which similarly can be described as a choice of $q \leq j$ non-members using resource k and a partition of the $j-q$ remaining non-members to the first $k-1$ resources (which corresponds to a strategy profile of a non-member agent set B' of $j-q$ agents in the $(k-1)$ -subgame). The safety level strategy for a coalition C in the k -subgame is therefore a composition of the safety level strategy for a coalition of $|C| - p$ in the $(k-1)$ -subgame and p agents using resource k for some $p \leq |C|$. \square

Thus, by Theorems 1 and 2, for SLC^{SRSG} -games we can efficiently test the core non-emptiness, construct a core imputation or, given one, check if it belongs to the core. These results do not extend to all $SLC^{(S)CG}$ -games. When agents are allowed multiple-resource strategies (even if these are derived from a common set), computing colitional values becomes hard.

Theorem 3. *Computing the value of a coalition in SLC^{SCG} -games is NP-hard.*

Proof. We reduce from Exact-Cover-By-3-Sets (X3C). Consider an X3C instance, with a set $S = \{1, 2, \dots, 3m\}$ of $3m$ elements, and triplets S_1, \dots, S_n where $S_i \subset S$ and $|S_i| = 3$. We are asked whether there is an exact cover of S that uses exactly m (disjoint) triplets. We construct an SLC^{SCG} -game, where each element $r \in S$ corresponds to a resource r (that is, S corresponds to \mathbf{R}), where each resource r 's utility function satisfies $u_r(1) = 1$; $u_r(k) = 0$ for $k \geq 2$. The SLC^{SCG} -game has $N = m$ agents, and an agent is allowed to choose any resource triplet, S_j , from the given collection of triplets—that is, the strategy space of any agent i is given by $\mathbf{S}_i = \{S_1, \dots, S_n\}$. This is an anonymous game. Let $v(\mathbf{N})$ be the value of the grand coalition. We show that if the X3C is a “yes” instance, $v(\mathbf{N}) = 3m$ and if it is a “no” instance then $v(\mathbf{N}) < 3m$. Suppose the X3C is a “yes” instance, and let $S_{i_1}, S_{i_2}, \dots, S_{i_m}$ be the triplets in the exact cover. Let agent x choose the resources in S_x (for $x \in \{1, \dots, m\}$). Since S is an exact cover, each resource r is selected exactly once, so $v(\mathbf{N}) = 3m$. On the other hand, if the X3C is a “no” instance, any choice of m (or more) triplets S_{i_1}, \dots, S_{i_m} results in choosing at least one of the resources, r , more than once. Thus, the congestion on this resource results in a utility of 0 for all agents using it, so $v(\mathbf{N}) < 3m$. \square

4.2 Non-anonymous settings

We now turn to consider general, asymmetric settings where agents may have different strategy spaces. First, we observe that though these settings are anonymous in the original—non-cooperative—context, their corresponding SLC-games are not such:

Lemma 5. *The SLC^{CG} -games are, in general, non-anonymous.*

Proof. To see this, consider a simple (though extreme) example, in which there exists an agent, say i , that has an exclusive right to use a special resource rewarding its user with a very high utility, H . Now, any coalition C that includes i can guarantee to itself the utility of at least H , regardless of what the rest of agents do, while any coalition $C \setminus \{i\} \cup \{j\}$, $j \in \mathbf{N} \setminus C$, will never achieve this value. \square

Next, we show that losing anonymity results in high complexity of computing safety level values even for “degenerate” coalitions consisting of only a single agent:

Theorem 4. *Computing values of singleton coalitions in SLC^{CG} -games is coNP-hard.*

Proof. We reduce from the dominating-set (DS) problem. In DS, we are given a graph $G = \langle V, E \rangle$, and have to decide if there is a dominating vertex set of size at most K . A dominating vertex set is a vertex set $V' \subset V$ such that for every $v \in V$, either $v \in V'$ or $(u, v) \in E$ for some $u \in V'$. Denote $|V| = m$. We create an SLC^{CG} -game instance as follows: The resources correspond to the vertices V , and we define one additional resource, r^* (so, $R = m + 1$). The congestion function for resources $r \in V$ is given by $u_r(1) = H$ where $H > 3m$ is a very high value; $u_r(k) = 0$ for $k \geq 2$, and for the “special” resource we have $u_{r^*}(k) = 2m - k + 1$ for $k = 1, \dots, m + 1$; $u_{r^*}(k) = 0$ for $k \geq m + 2$. For any vertex resource $r \in V$ we define an agent a_v , who can choose any single resource which is a neighbour of v or resource r^* , so $\mathbf{S}_{a_v} = \{\{u\} \mid u \in V, (u, v) \in E\} \cup \{r^*\}$. There is also additional agent, a^* , whose only strategy is to select all the resources together, so $\mathbf{S}_{a^*} = \{\mathbf{R} = \{v \mid v \in V\} \cup \{r^*\}\}$.

Since agent a^* literally has no choice but using its only strategy, the value it obtains only depends on the choices of the other agents. Note that if in a strategy profile s , a^* there exist a vertex resource r so that a^* is its only user, then a^* obtains a value of at least H from s . Thus, to minimize a^* 's utility, each of the vertex resources must be used by some other agent. Now, if there is no dominating set of size K ,

this requires more than K other agents, and hence, at most $m - K - 1$ other agents can use r^* in such a profile, so $v(\{a^*\}) \geq 2m - (m - K - 1) = m + K + 1$. If there is a K dominating set, the K outsiders can choose this dominating set, so a^* obtains a utility of 0 from the vertex resources, and having $m - K$ outsiders on r^* results on the utility of $2m - (m - K) = m + K$ for a^* from r^* , so $v(\{a^*\}) = m + K$. Thus, $v(\{a^*\}) > m + K$ if and only if the DS instance is a “no” instance. \square

Obviously, in the non-anonymous setting, different coalitions—even of the same size—may have different values. While by Theorems 3 and 4 computing the values of coalitions is hard, one may try to find a *maximal* value a coalition of size (at most) k , where $1 \leq k \leq N$, can achieve. However, this option is ruled out as well, by the following theorem:

Theorem 5. *Finding the value of the grand coalition in SLC^{CG} -games is NP-hard.*

Proof. We reduce from Maximum Satisfiability (MAX-SAT). In MAX-SAT, one is given a Boolean formula, and is asked to determine the maximum number of clauses that can be satisfied by any assignment. Given a MAX-SAT instance, we construct an SLC^{CG} -game as follows. There is a resource for each clause, and an agent for each variable. An agent for variable x can either choose all the clauses satisfied by x or all the clauses satisfied by $\neg x$ (thus choosing an assignment for variable x). The resource utility function is $u_r(k) = \frac{1}{k}$, $k = 1, \dots, N$, for each resource r . Thus, if k agents choose a clause, each of them gets the utility of $\frac{1}{k}$ from the clause, and all of them together get the total utility of 1 from that clause. Thereby, given a strategy profile, its total value to the grand coalition equals exactly the number of satisfied clauses. Hence, the maximal value the grand coalition can achieve is the MAX-SAT solution. \square

5 Conclusions

We defined a safety level cooperative game induced by a normal form game, and examined this concept on the class of congestion games. We showed that for symmetric resource selection games the safety level of a coalition can be computed in polynomial time, and so several problems related to the core, a core imputation, and the Shapley value are polynomial time decidable. However, computing coalitional safety levels of congestion game (even a symmetric one) with multiple-resource strategies is NP-hard. For non-symmetric games, even computing the value of singleton coalitions or the grand coalition is hard.

This suggests a number of open questions for future research. First, other solution concepts should be investigated in the context of safety level cooperative games. In particular, the least-core [11] and the nucleolus [16] (which exists even if the core is empty), and the Banzhaf index [4] (an alternative to the Shapley value), should be considered. Second, the application domain should be extended to other scenarios, rather than just those of congestion games. One key domain for such analysis are auctions of various kinds. Finally, perhaps the most important task is that of finding tractable classes of games, where the hardness results we have demonstrated for congestion games would not hold. In this direction, encouraged by the example of resource selection games, we first of all intend to examine the classes of games for which the computation of a Nash equilibrium can be done in polynomial time. Immediate candidates include the matroid congestion games [1] and congestion-averse games [5, 18]; we are also interested if the tractability of the computation of coalitional values in these classes depends on the anonymity property same way it did for RSGs.

References

- [1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6), 2008.
- [2] R. Aumann. Acceptable points in general cooperative n -person games. In *Contributions to the Theory of Games*, volume IV, pages 287–324. Princeton University Press, Princeton, NJ, 1959.

- [3] Y. Bachrach. Honor among thieves - collusion in multi-unit auctions. In *AAMAS-10*, page (to appear), 2010.
- [4] J. F. Banzhaf. Weighted voting doesn't work: a mathematical analysis. *Rutgers Law Review*, 19:317–343, 1965.
- [5] A. Byde, M. Polukarov, and N. R. Jennings. Games with congestion-averse utilities. In *SAGT-09*, pages 220–232, 2009.
- [6] C. Daskalakis and C. Papadimitriou. Computing equilibria in anonymous games. In *FOCS-07*, pages 83–93, 2007.
- [7] D. B. Gillies. *Some theorems on n-person games*. PhD thesis, Princeton University, 1953.
- [8] A. Hayrapetyan, E. Tardos, and T. Wexler. The effect of collusion in congestion games. In *STOC-06*, 2006.
- [9] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21:85–101, 1997.
- [10] S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *AAAI-05*, pages 489–494, 2005.
- [11] M. Maschler, B. Peleg, and L. S. Shapley. Geometric properties of the kernel, nucleolus, and related solution concepts. *Mathematics of Operations Research*, 4(4):303–338, 1979.
- [12] I. Milchtaich. Congestion models of competition. *American Naturalist*, 147(5):760–783, 1996.
- [13] R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [14] T. Roughgarden and E. Tardos. How bad is selfish routing. *Journal of the ACM*, 49(2):236–259, 2002.
- [15] O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *WINE-06*, 2006.
- [16] D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17(6):1163–1170, 1969.
- [17] L. S. Shapley. A value for n-person games. *Contrib. to the Theory of Games*, pages 31–40, 1953.
- [18] T. Voice, M. Polukarov, A. Byde, and N. R. Jennings. On the impact of strategy and utility structures on congestion-averse games. In *WINE-09*, pages 600–607, 2009.
- [19] Robert J. Weber. Probabilistic values for games. Cowles Foundation Discussion Papers 471R, Cowles Foundation, Yale University, 1977.
- [20] E. Winter. The Shapley Value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.
- [21] M. Yokoo, V. Conitzer, T. Sandholm, N. Ohta, and A. Iwasaki. Coalitional games in open anonymous environments. In *AAAI-05*, pages 509–514, 2005.

A Missing Proofs

Below are the proofs omitted from the main text.

Lemma 3 (Core of Anonymous Games). *Let v be an anonymous cooperative game over N agents \mathbf{N} , with a non-empty core. Denote $q = \frac{v(\mathbf{N})}{N}$. Then the symmetric payoff distribution (q, q, \dots, q) is an imputation in the core.*

Proof. An imputation $p = (p_1, \dots, p_N)$ must fulfil two requirements to be in the core: (i) $\sum_{i \in \mathbf{N}} p_i = v(\mathbf{N})$, and (ii) $\forall C, \sum_{i \in C} p_i \geq v(C)$. Consider a core imputation $p = (p_1, \dots, p_N)$. If $p_i = p_j$ for any $i, j \in \mathbf{N}$, then denote $q = p_i = p_j$ and due to the first requirement we have $p(\mathbf{N}) = \sum_{i \in \mathbf{N}} q = q \cdot \sum_{i \in \mathbf{N}} 1 = q \cdot N$ so $q = \frac{v(\mathbf{N})}{N}$ as required. Otherwise, take agents i and j whose payoffs are maximal and minimal, respectively, so $p_i > p_j$. Let $q' = \frac{p_i + p_j}{2}$ and $p' = (p_1, p_2, \dots, p'_i = q', \dots, p'_j = q', \dots, p_N)$ be a new imputation. Denote $d = \frac{p_i - p_j}{2}$, so $p_i = p_j + 2d$. The imputation p' is “more balanced” than p , as i ’s payment drops: $p'_i = p_i - d$, while j ’s payment rises: $p'_j = p_j + d$. We show that p' is also a core imputation. Any coalition C such that $i, j \notin C$ has equal total payments in p and p' : $p'(C) = p(C)$. Hence, if $p(C) \geq v(C)$ we also have $p'(C) \geq p(C)$ and the core conditions for C hold. Also note that for any coalition C such that both i and j are in C we also have $p'(C) = p(C)$ so all the core conditions hold. Now, since $p'_j > p_j$, for any coalition C with $i \notin C$ and $j \in C$ we have $p'(C) > p(C) \geq v(C)$. Finally, consider a coalition C with $i \in C$ and $j \notin C$: in p' , this coalition C gets a lower payoff of $p'(C) = p(C) - d$. Let $C_j = C \setminus \{i\} \cup \{j\}$. We have $p(C) = p(C_j) - p_i + p_j = p(C_j) - (p_j + 2d) + p_j = p(C_j) + 2d$. Thus, $p'(C) = p(C) - d = p(C_j) + d$. Now, since the core conditions hold for C_j , i.e. $p(C_j) \geq v(C_j)$, we have that $p'(C) \geq v(C_j) + d$. Since the game is anonymous, we have $v(C_j) = v(C)$ (the difference between C_j and C is the exchange of two agents), so $p'(C) \geq v(C) + d$. Hence, p' is a core imputation.

We can repeat the process of converting a core imputation to a “more balanced” core imputation. After at most N such steps, the maximal difference between any two payments, $\max_{i,j} |p_i - p_j|$ drops by half. Thus, after sufficient number of iterations and up to any required accuracy, we obtain the balanced imputation $p^* = (q, q, \dots, q)$ which is also a core imputation, as required. \square

Lemma 4 (Shapley Value of Anonymous Games). *Let v be an anonymous cooperative game over N agents \mathbf{N} . Denote $q = \frac{v(\mathbf{N})}{N}$. Then the Shapley value is the symmetric payoff distribution (q, q, \dots, q) . If the core exists, then the Shapley value is in the core.*

Proof. Due to the properties of the Shapley value [17, 20], the payoff for any two equivalent agents i, j in the Shapley value must be the same, so $\phi_i(v) = \phi_j(v)$.³ Since the game is anonymous, we have $\phi_i(v) = \phi_j(v)$ for any two agents, i, j . Thus, the Shapley value is $\phi(v) = (x, x, \dots, x)$ for some value x . Since by definition of the Shapley value we have $\sum_{i=1}^N x = v(\mathbf{N})$ we get that $x = q = \frac{v(\mathbf{N})}{N}$. Due to Lemma 3, if the core is non-empty, then $\phi(v) = (q, q, \dots, q)$ is in the core. \square

³In fact, this is one of the axioms characterizing the Shapley value.