

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

Yoram Bachrach<sup>1</sup>, Maria Polukarov<sup>2</sup>, and Nicholas R. Jennings<sup>2</sup>

<sup>1</sup> Microsoft Research, Cambridge, UK  
yobach@microsoft.com

<sup>2</sup> University of Southampton, UK  
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 interactions of selfish rational agents. An agent may not follow a “prescribed” behavior if deviating from it improves its utility, so *stable* outcomes are central in game theory. In *non-cooperative* games, where agents take individual actions, the prominent stability concept is the *Nash equilibrium*—a strategy profile where no agent has a beneficial unilateral deviation. However, it does not take into account collective deviations by groups of agents; The *strong equilibrium* [2]—a strategy profile with no profitable agent subset deviations—extends Nash equilibrium to coalitions. *Cooperative* games consider how coalitions of agents cooperate, focusing on how the utility is distributed among the agents. In a non-cooperative game, agents act independently based on their individual interests. In many scenarios traditionally modeled as non-cooperative games (e.g., auctions, network and congestion games), a coalition can *jointly* decide on a collective action and make monetary transfers to *share* the gains. This requires applying tools from cooperative game theory to such domains. Hayrapetyan et al [10] modeled coalitions in *congestion games*. In a congestion game [15], a set of agents shares a set of resources, and an agent’s strategy is to choose a subset of resources to use to minimize the sum of congestion-dependent costs over its selected resources. In [10], the authors assume that agents may collude to maximize their collective welfare. Their model allows monetary transfers but uses a different perspective

than cooperative game theory, focusing on the negative effect of collusion on the social welfare. Other papers examine coalition formation in multi-unit auctions [3], assuming non-colluders bid *truthfully* or fostering cooperation through *external subsidies* [4,14].

In contrast, we study *safety level* coalitions in competitive games. As opposed to standard cooperative games, the utility of a coalition 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 “punish” the coalition members and take actions that minimize the collaborators’ utility. 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 our model the collaborators are “good” to each other by coordinating actions and sharing gains and the non-collaborators are “bad” adversaries who reduce the collaborators’ utility. If the collaborators adopt this view of non-collaborators as adversaries, they must be “cautious” and prepare for the worst-case choice of the non-collaborators, using their joint safety level strategy. To do this, they can agree on monetary transfers through an enforceable contract for distributing the gains. Solution concepts such as the core [9] and the Shapley value [18] can be used to predict what transfers would occur. Several works [7,8,21,5] consider computing the core and Shapley value in various domains.

We examine safety level cooperation in congestion games. These games where self-ish agents choose from a common set of resources and derive individual utilities that depend on the total congestion on each resource, are fundamental to many applications [13,15,16]. Such games have Nash equilibria in pure strategies [15], and some restricted classes have strong equilibria in pure strategies [11,17]. Recent work focuses on specific subclasses that are computationally tractable [1,12]. One subclass which we also examine is *resource selection games*, where each agent chooses a *single* resource.

We 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 core 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 games<sup>1</sup>—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 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.

## 1.1 Preliminaries

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 reward an agent gets. 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

<sup>1</sup> Not to be confused with the *anonymity-proof* solutions [20] which are robust under “false name” manipulations. We refer to games where the characteristic function is not “sensitive” to the agents’ identities so equal size coalitions get equal values.

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:  $s_i^{SL} \in \arg \max_{s_i \in \mathbf{S}_i} \min_{s_{-i} \in \mathbf{S}_{-i}} U_i(s_i, s_{-i})$ .

Some utility functions ignore the identities of the agents, and only take into account *the number of times* each strategy is played. Settings where identities are irrelevant are *anonymous*. 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 maps 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 *symmetric games*, where the payoffs for playing a particular strategy are the same for different agents and depend only on the other strategies employed, so one can change the identities of the agents without changing the payoffs to the strategies. A game with strategy spaces  $\mathbf{S}_1 = \dots = \mathbf{S}_N = \mathbf{S}$  is symmetric if for any permutation  $\pi$  over  $\mathbf{N}$  and 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  $j = 1, \dots, N$ .

A *transferable utility cooperative game* has 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 all agents except  $i$  as  $\mathbf{N}_{-i} = \mathbf{N} \setminus \{i\}$ . A coalitional game is *monotone* if  $v(C') \leq v(C)$  for any  $C' \subseteq C$ .

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* [9]. 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* [18] which focuses on *fairness*. It fulfills several important fairness axioms [18] and has been used to fairly share gains or costs. It depends on the agent's marginal contribution to possible coalition permutations. We denote by  $\pi$  a permutation (ordering) of the agents, and by  $\Pi$  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  $m_{i_1}^\pi(v) = v(\{i_1\})$  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*, and contains the game's 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) [15], 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  $\mathbf{R} = \{r_1, \dots, r_R\}$  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:  $U_i(s) = \sum_{r \in s_i} u_r(h_r(s))$ . 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 resource is the same for each of its users, the utility any agent gets from 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 [15]. Resource selection games with monotone utility functions also admit strong equilibria [11]. 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.

## 2 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 agent set, and  $\mathbf{S}_i$  and  $U_i$  denote, respectively, strategy spaces and utility functions of *individual* agents. We are interested in scenarios where it makes sense to the agents to form coalitions and coordinate their actions to optimize their *collective* gains and take a *safety level* approach to analyzing gains of a coalition. 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. 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 coalition's utility 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. 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 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 its minimal utility 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 it can guarantee *as a whole* when its members cooperate. A key challenge is determining how the members would share this value. To answer this, we define a *safety level cooperative game* (SLC-game) for  $\Gamma$ :

**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  $\Gamma$ : for each  $C \subseteq \mathbf{N}$ ,  $v(C) = U_C^*$ .*

We write  $SLC^\Gamma$  to indicate that an SLC-game is induced by a game  $\Gamma$ . Regardless of their underlying games  $\Gamma$ , all SLC-games have the following 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 then that of a smaller one. Formally, denote  $B = \mathbf{N} \setminus C$ , so  $\mathbf{N} \setminus C' = B \cup D$ . We have:

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

## 2.1 Safety Level Coalitions in Congestion Games

We analyze safety level coalitions in congestion games and resource selection games. We 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 under monetary transfers. We show anonymous and non-anonymous SLC-games differ computationally.

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 game classes.

### 2.2 Anonymous Cooperative Games

We consider the properties of SLC-games induced by symmetric congestion games, where all agents use a common set of strategies. We start with Lemma 2 showing that these games satisfy *anonymity*. 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<sup>SCG</sup>-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  $C \cup \{j\}$ .

In anonymous games the Shapley value can be found in polynomial time and is in the core when it's not empty (proofs omitted for lack of space).

**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. Some safety level games have empty cores (see Examples 1 and 2 below). Empty cores 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.

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 are  $SLC^{CG}$ -games with non-empty core).* Consider 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  $r \in \mathbf{R}, k = 1, \dots, N$ . This 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.

*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). 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. They can either choose to both use the same resource, or to each use a different resource. 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.

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  $S_1 = S_2 = S_3 = \{\{a\}, \{b\}, \{c\}\}$ , 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 fulfill 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 a core imputation, otherwise the core is empty.

The computation of coalition values in SLC-games, can be difficult 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  $\epsilon$ , 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  $\epsilon$  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  $S_C$ , any strategy in  $S_{N \setminus C}$  of the remaining 2 agents results in all agents on  $a$  getting a utility of  $\epsilon$  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  $v(3) = 2 + \epsilon$ .

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

**Theorem 2.** For  $SLC^{SRSG}$ -games, computing safety level strategies 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  $S_i = \{r_1, r_2, \dots, r_k\} \subseteq \mathbf{R}$ . 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 use 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  $\sum_{j=1}^{k-1} c_x = i$ . 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  $\sum_{j=1}^{k-1} b_x = j$ . 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  $p \leq |C|$ .

By Theorems 1 and 2, for  $SLC^{SRSG}$ -games we can efficiently test core non-emptiness, construct a core imputation or check if an imputation is in the core. These results do not extend to all  $SLC^{(S)CG}$ -games. When agents are allowed multiple-resource strategies (even if derived from a common set), computing coalitional values is 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  $v(\mathbf{N}) < 3m$ .

### 2.3 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 single agent  $i$  that has an exclusive right to use a special resource rewarding its user with a very high utility,  $H$ . Any coalition  $C$  that includes  $i$  guarantees itself a utility of at least  $H$ , regardless of what the rest of agents do, while any  $C \setminus \{i\} \cup \{j\}$ ,  $j \in \mathbf{N} \setminus C$ , cannot achieve this value.

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 dominating-set (DS). 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  $K$ . In a dominating vertex set  $V' \subset V$  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 add a 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  $S_{a_v} = \{\{u\} \mid u \in V, (u, v) \in E\} \cup \{\{r^*\}\}$ . There is also additional agent,  $a^*$ , whose only strategy is to select all the resources, so  $S_{a^*} = \{\mathbf{R} = \{v \mid v \in V\} \cup \{r^*\}\}$ . Since  $a^*$  must use its only strategy, the value it obtains depends only on the choices of the other agents. If in a strategy profile  $s$ ,  $a^*$  there exists 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. If there is no dominating set of size  $K$ , this requires more than  $K$  other agents, so 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,  $v(\{a^*\}) > m + K$  iff the DS instance is a “no” instance.

In non-anonymous settings even equal size coalitions may have different values. While by Theorems 3 and 4 computing coalition values is hard, one may seek the *maximal* value of coalitions of size (at most)  $k$ , where  $1 \leq k \leq N$ . We show this is also hard.

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

*Proof.* We reduce from MAX-SAT, where given a Boolean formula we are asked to find the maximum number of clauses that can be satisfied by any assignment. Given a MAX-SAT instance, we construct a  $SLC^{CG}$ -game. There is a resource for each clause, and an agent for each variable. An agent for variable  $x$  can either choose all clauses satisfied by  $x$  or all 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 value to the grand coalition is exactly the number of satisfied clauses.

### 3 Conclusions

We defined a safety level cooperative game induced by a normal form game, and examined this concept on the class of congestion games. A number of questions remain open for future research. First, other solution concepts should be investigated in the context of safety level cooperative games. Second, the application domain should be extended to non-congestion scenarios, such as auctions. Finally, an important task is finding tractable game classes, where the our hardness results do not hold. We intend to examine games where computing Nash equilibria can be done in polynomial time, such as matroid congestion games [1] and congestion-averse games [6,19].

## References

1. Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55(6) (2008)
2. Aumann, R.: Acceptable points in general cooperative  $n$ -person games. In: *Cont. to the Theory of Games*, Princeton, NJ, vol. IV, pp. 287–324 (1959)
3. Bachrach, Y.: Honor among thieves - collusion in multi-unit auctions. In: *AAMAS* (2010)
4. Bachrach, Y., Elkind, E., Meir, R., Pasechnik, D., Zuckerman, M., Rothe, J., Rosenschein, J.: The cost of stability in coalitional games. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) *SAGT 2009*. LNCS, vol. 5814, pp. 122–134. Springer, Heidelberg (2009)
5. Bachrach, Y., Rosenschein, J.S.: Power in threshold network flow games. In: *Autonomous Agents and Multi-Agent Systems* (2009)
6. Byde, A., Polukarov, M., Jennings, N.R.: Games with congestion-averse utilities. In: Mavronicolas, M., Papadopoulou, V.G. (eds.) *SAGT 2009*. LNCS, vol. 5814, pp. 220–232. Springer, Heidelberg (2009)
7. Conitzer, V., Sandholm, T.: Computing Shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In: *AAAI* (2004)
8. Elkind, E., Goldberg, L.A., Goldberg, P., Wooldridge, M.: Computational complexity of weighted threshold games. In: *AAAI* (2007)
9. Gillies, D.B.: Some theorems on  $n$ -person games. PhD thesis, Princeton University (1953)
10. Hayrapetyan, A., Tardos, E., Wexler, T.: The effect of collusion in congestion games. In: *STOC 2006* (2006)
11. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. *Games and Economic Behavior* 21, 85–101 (1997)
12. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q.: Fast and compact: A simple class of congestion games. In: *AAAI 2005*, pp. 489–494 (2005)
13. Milchtaich, I.: Congestion models of competition. *American Naturalist* (1996)
14. Resnick, E., Bachrach, Y., Meir, R., Rosenschein, J.S.: The cost of stability in network flow games. In: Královíč, R., Nawiński, D. (eds.) *MFCS 2009*. LNCS, vol. 5734, pp. 636–650. Springer, Heidelberg (2009)
15. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67 (1973)
16. Roughgarden, T., Tardos, E.: How bad is selfish routing. *J. ACM* 49(2), 236–259 (2002)
17. Rozenfeld, O., Tennenholtz, M.: Strong and correlated strong equilibria in monotone congestion games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 74–86. Springer, Heidelberg (2006)
18. Shapley, L.S.: A value for  $n$ -person games. *Cont. to the Theory of Games* (1953)
19. Voice, T., Polukarov, M., Byde, A., Jennings, N.R.: On the impact of strategy and utility structures on congestion-averse games. In: Leonardi, S. (ed.) *WINE 2009*. LNCS, vol. 5929, pp. 600–607. Springer, Heidelberg (2009)
20. Yokoo, M., Conitzer, V., Sandholm, T., Ohta, N., Iwasaki, A.: Coalitional games in open anonymous environments. In: *AAAI 2005*, pp. 509–514 (2005)
21. Zuckerman, M., Faliszewski, P., Bachrach, Y., Elkind, E.: Manipulating the quota in weighted voting games. In: *AAAI* (2008)