Taxation and Stability in Cooperative Games

Yair Zick
School of Physical and Mathematical Sciences
Nanyang Technological University, Singapore
yair0001@ntu.edu.sg

Maria Polukarov
School of Electronics and Computer Science
University of Southampton, UK
mp3@ecs.soton.ac.uk

Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton, UK
nrj@ecs.soton.ac.uk

ABSTRACT

Cooperative games are a useful framework for modeling multi-agent behavior in environments where agents must collaborate in order to complete tasks. Having jointly completed a task and generated revenue, agents need to agree on some reasonable method of sharing their profits. One particularly appealing family of payoff divisions is the *core*, which consists of all coalitionally rational (or, stable) payoff divisions. Unfortunately, it is often the case that the core of a game is empty, i.e. there is no payoff scheme guaranteeing each group of agents a total payoff higher than what they can get on their own.

As stability is a highly attractive property, there have been various methods of achieving it proposed in the literature. One natural way of stabilizing a game is via taxation, i.e. reducing the value of some coalitions in order to decrease their bargaining power. Existing taxation methods include the ε-core, the least-core and several others.

However, taxing coalitions is in general undesirable: one would not wish to overly tamper with a given coalitional game, or overly tax the agents. Thus, in this work we study minimal taxation policies, i.e. those minimizing the amount of tax required in order to stabilize a given game. We show that games that minimize the total tax are to some extent a linear approximation of the original games, and explore their properties. We demonstrate connections between the minimal tax and the cost of stability, and characterize the types of games for which it is possible to obtain a tax-minimizing policy using variants of notion of the ε-core, as well as those for which it is possible to do so using reliability extensions.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms

Theory

Keywords

Cooperative Games, Core Stability, Optimal Taxation Schemes

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

1. INTRODUCTION

The theory of cooperative games with transferable utility (TU games) has been widely used in multi-agent systems to study scenarios where groups of agents may form coalitions and generate profits. Formally, a *TU cooperative game* $\mathcal{G}$ is given by a set of agents $N = \{1, \dots, n\}$ and a *characteristic function* $v : 2^N \to \mathbb{R}$, assigning a value to each subset $S \subseteq N$. It is often assumed that agents will form the *grand coalition*, i.e. the set of all agents N. However, in some cases agents may form *coalition structures* [1]; that is, agents split into disjoint coalitions, working independently in order to maximize total revenue. Our focus is on the former approach, where the grand coalition is formed.[1]

Having formed the grand coalition, agents must decide on some reasonable payoff division. Payoff division schemes are known in the literature as *solution concepts* (see [14] and [5] for a review of common solution concepts); a solution concept is a mapping whose input is a cooperative game $\mathcal{G}$, and whose output is a set of payoff divisions. The *core* [7] is arguably the most prominent solution concept; a payoff division is said to be in the core of $\mathcal{G}$ (denoted $Core(\mathcal{G})$) if no subset of agents can get more by leaving the group and working on their own, i.e. the total payoff to every subset of agents S is at least its value $v(S)$. Thereby, core outcomes capture the notion of *stability* in cooperative games; this is because under a core outcome, no subset of agents can profitably deviate. Unfortunately though, many cooperative games have an empty core; this means that no matter how agents divide their profits, there will always be a subset of agents that is paid less than what it can make on its own.

1.1 Relaxing the Core Requirement

Core stability is a highly desirable, but rarely achievable, property; thus, one would ideally like to maintain a set of "somewhat" stable payoffs when the core is empty. This can be done via several approaches. First, one may drop the stability requirement and focus on other types of solution concepts, for which a payoff division is guaranteed to exist. Such solutions, in particular, include the *nucleolus* [15] and the *bargaining set* [6], and have their own justifications and appeal: for example, the nucleolus of a cooperative game is a payoff division scheme that minimizes some measure of unhappiness in the game.

An alternative approach would be to stabilize the game

[1]We do not use coalition structures both for the sake of simplicity and in order to maintain consistency with other solution concepts, which often do not utilize coalition structures in their definitions.

via external subsidies. Intuitively, a game is not stable since the grand coalition is unable to generate enough revenue to satisfy the demands of each subset of agents. An external party that is interested in stabilizing the game provides a subsidy to the agents if they form the grand coalition, and thus a value of $\delta v(N)$ is divided among them, where $\delta \geq 1$. Clearly, any game can be stabilized using a large enough δ; however, the external party would naturally be interested in the *minimal subsidy* required in order to stabilize the game.

Finally, a game can be stabilized by relaxing the core constraints. For example, it is often reasonable to assume that a subset of agents would not choose to deviate if the additional payment they can secure by deviating does not exceed some $\varepsilon > 0$, i.e. a coalition will choose to deviate only if a substantial gain can be made by deviating, as deviation itself is a costly act. Alternatively, the ε can be thought of as a *tax* imposed on a coalition should it choose to deviate; again, this can be viewed as some external party wishing to stabilize the game, but doing so via reducing the bargaining power of subsets of N, rather than increasing the desirability of N itself. Formally, given a game $\mathcal{G}$, the game $\mathcal{G}_\varepsilon$ has the value of every coalition except N reduced by some ε. It is easy to see that for a large enough ε, the game $\mathcal{G}_\varepsilon$ is stable. Note that it can also be assumed that $\varepsilon < 0$; that is, if the game is stable, it may be possible to add a value of ε to the value of each coalition $S \subseteq N$ (except N itself), thus increasing the bargaining power of sub-coalitions. Since our focus in this work is on stabilizing games whose cores are empty, we assume from now on that $\varepsilon \geq 0$. Again, we are interested in the *smallest* possible ε for which $\mathcal{G}_\varepsilon$ is stable, as that minimal ε corresponds to the minimal change required in order to stabilize the game via ε reductions. This gives rise to the notion of the *strong least core*, which corresponds to the ε^*-core where ε^* is the smallest value for which the game $\mathcal{G}_{\varepsilon^*}$ has a non-empty core. Variants of the strong-ε core exist, each corresponding to a different method of taxation (see Section 2 for a more detailed discussion).

The study of taxation as a method of stabilizing cooperative games is not recent; these notions have been first explored in [9], where the geometry of the least core and its connection to the nucleolus and other solution concepts have been established. More recently, Bejan and Gómez [4] study the properties of individual taxation schemes. Specifically, given a game $\mathcal{G}$, they consider various ways in which an individual taxation scheme (i.e. a mapping from a game to a vector in $\mathbb{R}^n$) can stabilize a game, and explore their properties. Moreover, Bejan and Gómez show a connection between an optimal individual taxation scheme (i.e. one that minimizes the total amount of tax taken from individuals) and the cost of stability (i.e., the minimal subsidy to the value of the grand coalition required in order to stabilize the game). The goal of their work is to provide axioms which would hold for those taxation schemes that coincide with the core whenever the core is not empty, i.e. taxation schemes that do not tax individuals at all if the game is stable to begin with. In this sense, our approach is similar: the taxation notion we propose results in the core if this is not empty. However, unlike Bejan and Gómez, we study *group taxation schemes*, where the value of each coalition is reduced by a certain value, until the resulting game is stable. As we demonstrate, taking this approach results in significantly lower taxes for coalitions.

Also in this line of work, Gonzales and Grabisch [8] study a generalization of the extended core proposed by Bejan and Gómez [4], where a tax is only employed on coalitions of size at most k. Specifically, they look for games that minimize the difference between a group taxation scheme and an individual taxation scheme. The reasoning behind this methodology is simple: when a tax (or a payoff) is set upon sets rather than individuals, the sets of agents must bargain among themselves in order to agree on the way in which the tax should be divided. Thus, studying group taxation that minimizes the number of taxes on sets makes sense, as it minimizes the amount of complicated bargaining that the agents must undergo. In contrast, in our setting, we do not address the way in which taxes are distributed among sub-coalition members; rather, our results show that individual taxation does naturally arise in the setting that we propose. In the class of games that we study, the value of a blocking coalition is replaced by a value set by an additive vector, whose coordinates induce an individual tax.

Independent of the previous work, Bachrach et al. [2] study the *cost of stability* in coalitional games. The cost of stability is the minimal external subsidy to the grand coalition that is required in order to stabilize a given cooperative game. The authors [2] provide bounds on the amount of subsidy required for various classes of cooperative games. Additional bounds have been recently provided by Meir et al. [10] and Meir et al. [11]. In our work, we provide an explicit upper bound for the class of superadditive, anonymous games, and explore its connection to the bound presented in [2] for the same class. We show that for small enough coalitions (namely, those whose size is at most $\frac{n}{2}$), an optimal taxation policy is guaranteed to impose a lower tax than that used when computing the cost of stability; moreover, we show a relation between the amount of savings induced and the size of a coalition, with smaller coalitions guaranteed better tax reductions than larger coalitions. This result is particularly appealing, as it is often assumed that smaller coalitions are likelier to form than larger ones, thus ensuring a low tax on such coalitions is paramount.

As can be seen, the cost of stability, the least core, the concepts studied by Bejan & Gomez [4] and Gonzales & Grabisch [8] are aimed at finding the minimal amount of change required to stabilize a given cooperative game $\mathcal{G}$. We continue in this line of work, with the aim of finding the minimal amount of *taxation* required in order to achieve core stability.

1.2 Our Contribution

Against this background, we analyze scenarios where minimal taxation is employed in order to achieve stability in cooperative games. In particular, given a cooperative game $\mathcal{G}$, we examine the set of all *stable* cooperative games that are dominated by $\mathcal{G}$, i.e. their characteristic functions give a lower value to each coalition $S \subseteq N$. These games include all games that correspond to the ε-core of $\mathcal{G}$ and, in particular, the game corresponding to the least-core of $\mathcal{G}$, the notions described by Bejan and Gómez, and other notions such as graph restrictions [12] and reliability extensions [3].

We then proceed to look at the set of games on the efficient face of the polytope of stable games dominated by $\mathcal{G}$. We give a complete characterization of these games, which we term *maximal-stable games*; briefly, one can construct a maximal-stable game by replacing the value of every blocking coalition (i.e. one violating the core constraints) by a value induced by an additive taxation scheme. We explore the connections between the minimal possible tax required in order to stabilize a cooperative game and *the cost of stability*, both additive and relative. In particular, for the class of superadditive, anonymous cooperative games we provide an explicit formulation of a taxation policy, as well as a guaranteed improvement as opposed to the individual taxation scheme induced by the cost of stability. Finally, we characterize the types of games for which the cooperative game induced by the least-core is maximal-stable, as well as games whose *reliability extension* [3] is maximal-stable.

1.2.1 Organization

The paper unfolds as follows. Section 2 contains necessary definitions and preliminary results. In Section 3, we define and characterize maximal-stable games. We discuss their relations with the cost of stability in Section 4, and compare worst-case optimal taxation with that used in [2] for anonymous games in Section 5. We then describe classes of games for which the taxation schemes induced by variants of the ε-core (Section 6) and the reliability extension (Section 7) are optimal. Finally, Section 8 concludes with directions for future work.

2. PRELIMINARIES

In this section, we provide necessary notation and define concepts needed for presentation of our results in the following sections.

2.1 Basic Definitions

A *cooperative game with transferable utility* (TU game) is a tuple $\mathcal{G} = \langle N, v\rangle$, where $N = \{1, \dots, n\}$ is the set of agents and $v : 2^n \to \mathbb{R}_+$ is the *characteristic function* of $\mathcal{G}$, which assigns a value to each coalition $S \subseteq N$. It is assumed that $\mathcal{G}$ is non-negative, i.e. $v(S) \geq 0$ for all $S \subseteq N$, and that $v(\emptyset) = 0$. Alternatively, a cooperative game can be viewed as a vector in $\mathbb{R}_+^{2^n}$, where the value of the set S corresponds to the coordinate whose binary encoding has 1 in its i-th bit if $i \in S$, and 0 otherwise.

A cooperative game $\mathcal{G}$ is called *monotone* if for all $S \subseteq T \subseteq N$ we have that $v(S) \leq v(T)$; it is *superadditive* if for all $S, T \subseteq N$ such that $S \cap T = \emptyset$ we have that $v(S) + v(T) \leq v(S \cup T)$, and *additive* if the latter holds with equality. We say that a game $\mathcal{G}$ is *anonymous* if the value of a coalition does not depend on the identity of its member agents; i.e. if there exists some $f : N \to \mathbb{R}$ such that for all $S \subseteq N$ we have that $v(S) = f(|S|)$.

Throughout the paper, vectors are written as boldface lowercase letters and sets of agents are written as uppercase letters. Given a vector $\mathbf{x} \in \mathbb{R}^n$ and a set $S \subseteq N$, we write $x(S) = \sum_{i \in S} x^i$. A *payoff division* is a vector $\mathbf{x} \in \mathbb{R}_+^n$, where x^i is the payment to agent i. We say that $\mathbf{x}$ is *efficient* if $x(N) = v(N)$, i.e. the total payoff to all agents equals the value of the grand coalition N. We say that $\mathbf{x}$ is *individually rational* if for all $i \in N$, $x^i \geq v(\{i\})$, i.e. if every agent receives at least what he can make on his own. A vector $\mathbf{x} \in \mathbb{R}^n$ is called an *imputation* if it is both efficient and individually rational. We denote the set of imputations over $\mathcal{G}$ by $Imp(\mathcal{G})$.

The *core* of a cooperative game $\mathcal{G} = \langle N, v\rangle$ is the set of all imputations $\mathbf{x} \in Imp(\mathcal{G})$ such that for all $S \subseteq N$ it holds that $x(S) \geq v(S)$. We denote the core of $\mathcal{G}$ by $Core(\mathcal{G})$. As mentioned in Section 1 and as the following example illustrates, it is possible that the core is empty: $Core(\mathcal{G}) = \emptyset$.

EXAMPLE 1. *Consider the* three-majority game *where $N = \{1, 2, 3\}$ and the value of every $S \subseteq N$ is 1 if $|S| \geq 2$ and is 0 otherwise. In this case, we have that $Core(\mathcal{G}) = \emptyset$. Indeed, consider an imputation $\mathbf{x} \in Imp(\mathcal{G})$. Since $x(N) = v(N) = 1$, there is some agent $i \in N$ whose payoff is greater than zero. This means that $x(N \setminus \{i\}) = 1 - x^i < 1$; however, $|N \setminus \{i\}| = 2$, so $v(N \setminus \{i\}) = 1$, and hence $\mathbf{x} \notin Core(\mathcal{G})$. We observe that the three-majority game is superadditive, thus superadditivity is not a sufficient (nor, indeed, necessary) condition for core non-emptyness.*

If $Core(\mathcal{G}) \neq \emptyset$ then there is a way to distribute the profits among the agents so that no subset of them has an incentive to deviate from the grand coalition and work on their own. Thus, we refer to games with a non-empty core as *stable*.

2.2 Dominated Games

For each game $\mathcal{G}$, we are interested in the particular set of stable games it defines, as follows.

DEFINITION 2. *Given a game $\mathcal{G} = \langle N, v\rangle$ and a game $\mathcal{G}' = \langle N, v'\rangle$, we say that $\mathcal{G}$* dominates *$\mathcal{G}'$ if for all $S \subseteq N$ we have $v(S) \geq v'(S)$; if $\mathcal{G}$ dominates $\mathcal{G}'$, we write $\mathcal{G}' \leq \mathcal{G}$. The* dominated stable set *of $\mathcal{G}$, $\mathcal{S}(\mathcal{G})$, is the set of all stable games dominated by $\mathcal{G}$; that is,*

$$\mathcal{S}(\mathcal{G}) = \{\mathcal{G}' \in \mathbb{R}^{2^n} \mid \mathcal{G}' \leq \mathcal{G}, Core(\mathcal{G}') \neq \emptyset\}.$$

The set $\mathcal{S}(\mathcal{G})$ can be viewed as the result of the following process. Suppose a given game $\mathcal{G}$ is unstable; however, the game can be stabilized by reducing the value of certain coalitions. For example, a central authority may reduce the value of certain coalitions via taxation or fines; alternatively, once the grand coalition is formed, other (unformed) sub-coalitions suffer an erosion in their value as time goes by.

Some common solution concepts can be viewed as taxation schemes; we now describe those relevant to our work.

2.2.1 The Least Core and Relative Least Core

The ε*-core* [9] is an example of taxation where all coalitions suffer equal depreciation, i.e. a uniform value of ε is removed from each coalition. We note that while the value of each coalition is depreciated by ε, the total payoff divided is still $v(N)$, which motivates the following definition. Given a cooperative game $\mathcal{G} = \langle N, v\rangle$ and an $\varepsilon > 0$, the game $\mathcal{G}_\varepsilon = \langle N, v_\varepsilon\rangle$ is defined as follows: $v_\varepsilon(N) = v(N)$, but $v_\varepsilon(S) = v(S) - \varepsilon$ for all $S \subsetneq N$. Clearly, there is some $\varepsilon > 0$ for which $Core(\mathcal{G}_\varepsilon) \neq \emptyset$; the game induced by the *least core* [9] corresponds to the smallest such ε, denoted ε^*. One can similarly define the *relative ε-core* using $v_{rel(\varepsilon)}(S) = \varepsilon v(S)$, i.e. instead of taxing coalitions via lump-sum taxation, we take a proportional share of every coalition's value. Finally, the weak ε-core is defined via $v_{w(\varepsilon)} = v(S) - |S|\varepsilon$. That is, each agent $i \in N$ is charged a value of ε should he choose to join a deviating coalition.

2.2.2 Reliability Extensions

Given a game $\mathcal{G} = \langle N, v\rangle$ and a vector $\mathbf{r} \in [0, 1]^n$, the *reliability extension* of $\mathcal{G}$ by $\mathbf{r}$ is given by $\mathcal{G}_\mathbf{r} = \langle N, v_\mathbf{r}\rangle$, where $v_\mathbf{r}(S) = \sum_{C \subseteq S} v(C) \prod_{j \in C} r^j \prod_{k \in S \setminus C}(1 - r^k)$. The idea behind these games is that given a cooperative game, every agent $i \in N$ may fail with probability $1 - r^i$, and survive with probability r^i. Thus, $v_{\mathbf{r}}(S)$ is the expected value of the coalition S. These games have been studied in [3], and are similar in nature to multilinear extensions [13]. We note that if we assume that $\mathcal{G}$ is monotone, then $v_{\mathbf{r}}(S) \leq v(S)$ for all $S \subseteq N$; moreover, one can apply the reliability extension only to strict subsets of N, leaving the value of N under $v_{\mathbf{r}}$ as $v(N)$. This makes reliability extensions consistent with the notion of taxation: for each agent $i \in N$, there is some probability r^i that he will be removed from the game if he decides to join a deviating coalition. This makes the value of a deviating coalition S simply its expected value given the probability of agents in S being removed from the game.

3. MAXIMAL-STABLE GAMES

In this section, we analyze what we term *maximal-stable* games, which are defined as maximal elements of the dominated stable set, under the dominance relation.

DEFINITION 3. *A game $\mathcal{G}^*$ in the dominated stable set $\mathcal{S}(\mathcal{G})$ is called* maximal-stable *if there is no game $\mathcal{G}' \in \mathcal{S}(\mathcal{G})$ such that $\mathcal{G}^* \leq \mathcal{G}'$ and $\mathcal{G}' \neq \mathcal{G}^*$.*

Maximal-stable games are those games for which no coalition's value can be increased without either breaking the dominance condition (i.e. the value of the coalition S under $\mathcal{G}^*$ equals the value of S under $\mathcal{G}$), or the stability condition (i.e., for any $\varepsilon > 0$, there is no way to stabilize the game induced by increasing the value of S under $\mathcal{G}^*$ by ε). These are games for which the total reduction in coalitions' value is minimal. In a sense, they represent *locally optimal* taxation schemes; they reduce the value of coalitions in the minimal possible manner, while preserving stability. As we show below, such optimal taxation schemes will never reduce the value of the grand coalition.

PROPOSITION 4. *If $\mathcal{G}^* = \langle N, v^* \rangle$ is a maximal-stable game in $\mathcal{S}(\mathcal{G})$, then $v^*(N) = v(N)$.*

PROOF. Suppose by the contrary that $v^*(N) < v(N)$. Since $Core(\mathcal{G}^*) \neq \emptyset$, there is some $\mathbf{x} \in Core(\mathcal{G}^*)$. Let $\mathcal{G}' = \langle N, v' \rangle$ be defined as follows: $\forall S \subsetneq N$, $v'(S) = v^*(S)$, and $v'(N) = v(N)$. Moreover, let $\mathbf{y} \in \mathbb{R}^n$ be defined as follows:

$$y^i = x^i + \frac{v(N) - v^*(N)}{n}.$$

Since $y^i > x^i$ for all $i \in N$, we have that $y(S) > x(S) \geq v^*(S) = v'(S)$ for all $S \subsetneq N$. Furthermore, $y(N) = x(N) + v(N) - v^*(N) = v(N)$, and thus $\mathbf{y}$ is in the core of $\mathcal{G}'$. We conclude that $Core(\mathcal{G}') \neq \emptyset$, $\mathcal{G}^* \leq \mathcal{G}' \leq \mathcal{G}$, and $\mathcal{G}' \neq \mathcal{G}^*$; hence, $\mathcal{G}^*$ is not maximal-stable—a contradiction. □

Proposition 4 asserts that in any maximal-stable game the value of the grand coalition need not be reduced; essentially, this means that while the bargaining power of the agents is reduced, their social welfare is not: their total payoff is the same as what they would have received under $\mathcal{G}$. Furthermore, as the following propositions demonstrate, maximal-stable games are in some sense an additive approximation of the original game $\mathcal{G}$.

PROPOSITION 5. *If $\mathcal{G}^*$ is maximal-stable, then for all $S \subseteq N$ and all $\mathbf{x} \in Core(\mathcal{G}^*)$, $x(S) > v^*(S) \Leftrightarrow v^*(S) = v(S)$.*

PROOF. We know that if $\mathbf{x} \in Core(\mathcal{G}^*)$ then $x(S) \geq v^*(S)$ for all $S \subseteq N$. Now, if there is some $S' \subseteq N$ such that $x(S') > v^*(S')$ but $x(S') \leq v(S')$, we can create the game $\mathcal{G}' = \langle N, v' \rangle$ such that $v'(S) = v^*(S)$ for all $S \neq S'$ and $v'(S') = x(S')$. Note that $\mathcal{G}'$ is stable (in particular, $\mathbf{x} \in Core(\mathcal{G}')$), which contradicts the fact that $\mathcal{G}^*$ is maximal-stable. Finally, if $x(S') > v(S') > v^*(S')$, we can define the game $\mathcal{G}'$ to be $v'(S) = v^*(S)$ for all $S \neq S'$, and $v'(S') = v(S')$ for S'. Again, $\mathcal{G}' \geq \mathcal{G}^*$ and $\mathcal{G}' \neq \mathcal{G}^*$, and $\mathcal{G}'$ is stable ($\mathbf{x} \in Core(\mathcal{G}')$)—a contradiction. □

We use Proposition 5 to prove the following claim.

PROPOSITION 6. *For any $\mathbf{x}, \mathbf{y} \in Core(\mathcal{G}^*)$, if for some $S \subseteq N$ it holds that $x(S) < v(S)$, then $y(S) = x(S) = v^*(S)$.*

PROOF. Since $x(S) < v(S)$, by Proposition 5, $x(S) = v^*(S)$; if $y(S) \leq v(S)$ then $y(S) = v^*(S)$ as well, and we are done. Now, suppose $y(S) > v(S)$; since $x(S) = v^*(S) < v(S)$, then again by Proposition 5, it is impossible that $y(S) > v(S) > v^*(S)$. Thus, it must be that $y(S) = v^*(S) = x(S)$ This completes the proof. □

Propositions 5 and 6 imply that a maximal-stable game $\mathcal{G}^* = \langle N, v^* \rangle$ is defined by the original game $\mathcal{G}$ and some imputation $\mathbf{x} \in Imp(\mathcal{G})$, such that $v^*(S) = \min\{v(S), x(S)\}$ for all $S \subset N$. Some subsets of N are allowed to keep their original bargaining power, i.e. $v^*(S) = v(S)$, whereas some subsets have their bargaining power reduced via some additive measure $\mathbf{x}$ so that $v^*(S) = x(S) < v(S)$. This can be thought of as follows: agents are paid according to some $\mathbf{x} \in Imp(\mathcal{G})$; if $\mathbf{x} \notin Core(\mathcal{G})$ then blocking coalitions have their value set to $x(S)$, which entails stability. A game is maximal-stable if no coalition had its value unjustly reduced. Indeed, one can view a maximal-stable game as an additive game, but with some of the strict subsets of N having reduced values. Note that Proposition 6 does not imply that the vector $\mathbf{x}$ is unique; for that to hold, Proposition 6 must hold on all singleton coalitions, which may well not be the case. That is, despite being very close to being additive, the core of a maximal-stable game is not necessarily a singleton, as the following example illustrates.

EXAMPLE 7. *Consider a game $\mathcal{G} = \langle N, v \rangle$ defined with $N = \{1, \ldots, 4\}$, $v(\{1,2\}) = 1$, $v(\{3,4\}) = 4$, $v(N) = 4$ and $v(S) = 0$ for the rest of the coalitions. It is easy to see that $\mathcal{G}$ is not stable. However, reducing the value of $v(\{3,4\})$ to 3 will stabilize the game. Under the new game, core outcomes are all vectors of the form $\mathbf{x} = (x_1, x_2, x_3, x_4)$ where $x_1 + x_2 = 1$ and $x_3 + x_4 = 3$.*

However, Proposition 6 does imply that if the imputation generating the maximal-stable game is not unique, then, intuitively, *individual* bargaining power is low: for all $i \in N$ the values of the singleton coalitions need not to be changed. In such scenarios, individuals have their bargaining power reduced by taxation when they form coalitions, but the actual tax they incur does not come into play and does not affect their individual bargaining power. We do mention that for some games (e.g. strictly subadditive games), individual bargaining power may be reduced due to low non-singleton coalition values.

We observe that while there may be different individual taxation schemes that lead to identical maximal-stable games, the cores that they induce never coincide. Furthermore, an immediate corollary of Proposition 6 is that the cores of different maximal-stable games are disjoint.

PROPOSITION 8. *Given two maximal-stable games $\mathcal{G}^*$ and $\mathcal{G}'$ over $\mathcal{G}$, either $Core(\mathcal{G}^*) \cap Core(\mathcal{G}') = \emptyset$ or $\mathcal{G}' = \mathcal{G}^*$.*

PROOF. Assume there is some $\mathbf{x} \in Core(\mathcal{G}^*) \setminus Core(\mathcal{G}')$. Since $\mathbf{x} \in Core(\mathcal{G}^*)$ but is not in the core of $\mathcal{G}'$, there is some set $T \subseteq N$ such that $v^*(T) \leq x(T) < v'(T)$; since $v'(T) \leq v(T)$, it must be the case that $x(T) < v(T)$. By Proposition 6, for all $\mathbf{y} \in Core(\mathcal{G}^*)$ it must hold that $x(T) = y(T)$, which implies that for all $\mathbf{y} \in Core(\mathcal{G}^*)$ we get $y(T) < v'(T)$; hence, we have that $\mathbf{y} \notin Core(\mathcal{G}')$ for all $\mathbf{y} \in Core(\mathcal{G}^*)$.

Now, suppose that $Core(\mathcal{G}') = Core(\mathcal{G}^*)$. We write $\mathcal{C} = Core(\mathcal{G}') = Core(\mathcal{G}^*)$. Consider a set $S \subseteq N$; if there is some $\mathbf{x} \in \mathcal{C}$ such that $x(S) > v'(S)$ then $x(S) > v^*(S)$ as well, and we have that $v'(S) = v^*(S) = v(S)$. A similar argument shows that if there is some $\mathbf{x} \in \mathcal{C}$ such that $x(S) > v^*(S)$ then $v'(S) = v^*(S) = v(S)$. Now, if for all $\mathbf{x} \in \mathcal{C}$ neither of the above holds, then trivially we have $v'(S) = v^*(S) = x(S)$. This completes the proof. □

Finally, we note that a game $\mathcal{G}$ indeed may induce several maximal-stable games.

EXAMPLE 9. *Consider again the 3-majority game. It is easy to check that reducing the value of any coalition of size 2 to 0 generates a maximal-stable game. Also, setting the value of all coalitions of size 2 to $\frac{2}{3}$ generates a maximal-stable game. In the former case, let $i \in \{1, 2, 3\}$ be the agent that is in both coalitions of size 2 that retain a value of 1; the imputation $\mathbf{y}$ where $y^j = 0$ if $j \neq i$ and $y^i = 1$ is the one inducing the maximal stable game. For the latter case, the imputation $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ induces the maximal-stable game.*

Using these observations, we now formulate a linear program to find maximal-stable games of a given game $\mathcal{G}$.

3.1 LP formulation

Maximal-stable games reduce the values of coalitions in order to achieve stability; however, there are several possible maximal-stable games, and it is possible that some games achieve stability by taxing more than others. We are interested in finding maximal-stable games whose total reduction is minimal (or, in other words, those that generate maximal overall coalitional value). That is, we wish to find maximal-stable games that do not differ by much from the original game. One possible measure is finding maximal-stable games whose distance from the original game is minimal (as measured according to some norm). Using the maximum-norm as our measure, we obtain the following LP:

$$\begin{aligned} \text{max:} \quad & \textstyle\sum_{S \subseteq N} U_S && (1)\\ \text{s.t.} \quad & U_S \leq v(S) && \forall S \subseteq N \\ & U_S - x(S) \leq 0 && \forall S \subseteq N \\ & x(N) = v(N) && \\ & x^i \geq 0 && \forall i \in N \end{aligned}$$

In LP (1) above, the first set of constraints are *dominance constraints*, the second are *stability constraints*, and the last is an *efficiency constraint*. Note that we can assume with no loss of generality that U_S are unconstrained: the dominance and stability constraints simply state that $U_S \leq \min\{x(S), v(S)\}$, both non-negative values. If there is some set $T \subseteq N$ for which $U_T < 0$, we can set $U_T = \min\{x(T), v(T)\}$ and get a better solution; thus in any optimal solution all U_S are non-negative.

We first demonstrate that solutions to the above optimization problem define maximal-stable games.

PROPOSITION 10. *Let $((U^*_S), \mathbf{x}^*)$ be an optimal solution to LP 1. Then, the game $\mathcal{G}^* = \langle N, v^* \rangle$ where $v^*(S) = U^*_S$ for all $S \subseteq N$, is maximal-stable.*

PROOF. Given an optimal solution $((U^*_S), \mathbf{x}^*)$ to LP (1), suppose that the game $\mathcal{G}^*$ defined by $v^*(S) = U^*_S$, $S \subseteq N$, is not maximal-stable. In particular, this means that there exists some coalition T and some $\varepsilon > 0$ such that increasing the value of T by ε results in a stable game dominated by $\mathcal{G}$. This game, along with one of its core imputations, defines a feasible solution to LP 1, thus contrudicting the assumption that $((U^*_S), \mathbf{x}^*)$ is an optimal solution. □

4. MAXIMAL-STABLE GAMES AND THE COST OF STABILITY

As previously mentioned, another common approach to stabilize a cooperative game is by increasing the desirability of the grand coalition via external subsidies, rather than reducing the bargaining power of its subsets by taxation. Formally, given a game $\mathcal{G} = \langle N, v \rangle$, one may subsidize the grand coalition up to the value of $\bar{v}(N) = v(N) + \Delta$ for $\Delta \geq 0$ (or, $\bar{v}(N) = \delta v(N)$ for $\delta \geq 1$) so that the resulting game $\bar{\mathcal{G}} = \langle N, \bar{v} \rangle$, where $\bar{v}(S) = v(S)$ for all $S \subsetneq N$, is stable. The smallest value of Δ (respectively, δ) for which the stability is achieved, defines the additive (respectively, relative) cost of stability.

It was shown by Bejan and Gómez [4], that the additive cost of stability equals the total tax required in order to stabilize a game, when one implements a taxation policy which taxes individuals rather than sets. We now turn to consider the relation between the cost of stability and maximal stability in the general setting with coalitional taxation schemes, as defined in Section 3. To this end, we utilize LP duality in order to draw similarities between the two problems.

Let us first rewrite LP (1) as an equivalent minimization problem for finding an optimal taxation scheme, i.e. the minimal amount that needs to be deducted from each coalitional value as per $\mathcal{G}$. We just set $t_S = v(S) - U_S$ and get:

$$\begin{aligned} \text{min:} \quad & \textstyle\sum_{S \subseteq N} t_S && (2)\\ \text{s.t.} \quad & x(S) + t_S \geq v(S) && \forall S \subseteq N \\ & x(N) = v(N) && \\ & t_S, x^i \geq 0 && \forall S \subseteq N, \forall i \in N \end{aligned}$$

Bejan and Gómez [4] employ individual, rather than coalitional, taxation: that is, in their work, a taxation scheme is given by a vector $\mathbf{t} = (t^1, \ldots, t^n)$, and they seek to minimize the total sum of individual taxes, $t(N) = \sum_{i \in S} t^i$. The tax on a coalition S is simply the sum of individual taxes over the members of the coalition: that is, $t_S = t(S) = \sum_{i \in S} t^i$, except for the grand coalition on which the authors impose zero tax. The important result of [4] is that if $\mathbf{t}^*$ is a minimal taxation scheme, then the cost of stability equals $t^*(N)$.

Our setting is strictly more general than that of Bejan and Gómez [4]: we employ combinatorial rather than additive taxes, i.e. t_S need not to be equal to $t(S)$. Furthermore, our objective function is also different: we are interested in minimizing the total amount of tax from all coalitions, rather than from individuals. Recall from Proposition 4 that $t_N = 0$, and so $\sum_{S \subseteq N} t_S = \sum_{S \subsetneq N} t_S$. This finally implies that the optimal value of LP (2) is bounded from above by $\sum_{S\subseteq N} t(S) = \sum_{i=1}^{n} \sum_{S\neq N: i\in S} t^i = (2^{n-1}-1)t(N)$, for any feasible individual taxation scheme $\mathbf{t}$. Thus, given the result in [4], the value of our optimal taxation scheme is at most $(2^{n-1}-1)$ times the cost of stability, which, in turn, is given by the following LP:

$$\begin{aligned} \text{min:} \quad & (2^{n-1}-1)t(N) \\ \text{s.t.} \quad & x(S)+t(S) \geq v(S) \quad \forall S \subseteq N \\ & x(N) = v(N) \\ & t^i, x^i \geq 0 \qquad \forall i \in N \end{aligned} \tag{3}$$

Let us now observe the duals of the linear programs above. For LP (2) we get the following problem:

$$\begin{aligned} \text{max:} \quad & \sum_{S\subseteq N} \beta_S v(S) - \gamma v(N) \\ \text{s.t.} \quad & 0 \leq \beta_S \leq 1 \qquad \forall S \subseteq N \\ & \textstyle\sum_{S:i\in S} \beta_S \leq \gamma \qquad \forall i \in N \end{aligned} \tag{4}$$

and the dual of LP (3) is given by:

$$\begin{aligned} \text{max:} \quad & \sum_{S\subseteq N} \beta_S v(S) - \gamma v(N) \\ \text{s.t.} \quad & \textstyle\sum_{S:i\in S} \beta_S \leq 2^{n-1}-1 \quad \forall i \in N \\ & \textstyle\sum_{S:i\in S} \beta_S \leq \gamma \qquad \forall i \in N \\ & \beta_S \geq 0 \qquad \forall S \subseteq N \end{aligned} \tag{5}$$

Thus, instead of requiring that each β_S will be at most 1, as in LP (4), only a general bound on the total weight of the variables β_S is required in LP (5). In particular, if LP (5) has an optimal solution where all β_S are at most 1, then the cost of stability (normalized by $2^{n-1}-1$) equals the amount deducted by an optimal taxation policy, i.e. an individual taxation scheme will also be an optimal taxation scheme.

5. OPTIMAL TAXATION IN ANONYMOUS GAMES

In this section, we compute optimal taxation policies for the special case of anonymous cooperative games; furthermore, we show that in the case of superadditive, anonymous games, small coalitions enjoy a significant reduction in their taxation under an optimal taxation policy.

Recall that a cooperative game $\mathcal{G} = \langle N, v\rangle$ is anonymous if there exists some function $f : N \to \mathbb{R}$ such that for all $S \subseteq N$ we have that $v(S) = f(|S|)$. As we have previously shown, an optimal taxation policy is completely defined by an imputation vector $\mathbf{x}$, and the maximal-stable game it induces has $v^*(S) = \min\{x(S), v(S)\}$. Given a vector $\mathbf{x} \in \mathbb{R}^n$, let $\mathbf{x}_{ij}$ be the vector $\mathbf{x}$ with the i-th and j-th coordinates swapped. We first make the following observation:

Observation 11. *Given an anonymous game $\mathcal{G} = \langle N, v\rangle$, if $\mathbf{x}$ induces a maximal-stable game, then for any $i, j \in N$, $\mathbf{x}_{ij}$ induces a maximal-stable game.*

Now, observe the set of imputations corresponding to optimal solutions of LP (2); this is a convex set [16], thus if $\mathbf{x}, \mathbf{y}$ induce optimal taxation policies, so does any linear combination of them. In particular, this implies the following:

Lemma 12. *If $\mathcal{G}$ is an anonymous cooperative game with $v(S) = f(|S|)$, then the imputation vector $\mathbf{e}$ where $e^i = \frac{f(n)}{n}$ for all $i \in N$, induces an optimal taxation policy.*

Proof. To see why this is the case, take a vector $\mathbf{z}$ such that the function $D(\mathbf{x}) = \sum_{i=1}^{n}\sum_{j=1}^{n} |x^i - x^j|$ is minimized over the set of vectors inducing optimal solutions to LP (2). Suppose that there are some $i^*, j^* \in N$ such that $z^{i^*} > z^{j^*}$. Observe the vector $\mathbf{y} = \frac{\mathbf{z}+\mathbf{z}_{i^*j^*}}{2}$. Since the set of optimal solution s is convex, and since $\mathbf{z}_{i^*j^*}$ induces an optimal solution as well, $\mathbf{y}$ induces an optimal solution. However, we have that $D(\mathbf{z}) - D(\mathbf{y}) = 2|z^{i^*} - z^{j^*}| > 0$, hence $\mathbf{z}$ does not minimize the function $D(\mathbf{x})$, a contradiction. □

Using Lemma 12, we can now provide an explicit formulation for an optimal taxation policy for anonymous games.

Theorem 13. *Let $\mathcal{G}$ be an anonymous cooperative game with $v(S) = f(|S|)$. Then, taxing each set $S \subseteq N$ of size s by $t_S = \max\{0, f(s) - \frac{s}{n}f(n)\}$ is optimal.*

Proof. The theorem follows directly from Lemma 12 and Proposition 5, stating that the value of any set $S \subseteq N$ in a maximal-stable game is given by $v^*(S) = \min\{v(S), x(S)\}$. Hence, in the anonymous setting, the tax for a set S of size s is $\max\{0, v(S) - x(S)\} = \max\{0, f(s) - \frac{s}{n}f(n)\}$. □

We term the taxation policy presented in Theorem 13 above the *anonymous* taxation policy, as it does not differentiate between coalitions of the same size.

Now, as has been observed by Bachrach et al. [2], if $\mathcal{G}$ is anonymous and *superadditive*, then $f(s) \leq \frac{s}{n-s}f(n)$. Hence, in the optimal taxation policy, the most that a coalition is taxed is $(\frac{s}{n-s} - \frac{s}{n})f(n)$, which equals $\frac{s^2}{n(n-s)}f(n)$. The authors [2] also provide a (tight) bound on the cost of stability that is obtained by adding $\frac{f(n)}{n}$ to each agent's payoff. Thus, the total extra amount that is paid to each coalition of size s is $\frac{s}{n}f(n)$. Note that $\frac{s^2}{n(n-s)} \leq \frac{s}{n}$ if and only if $s \leq \frac{n}{2}$, so large coalitions may be taxed a higher amount under the optimal taxation policy; to conclude, for coalitions of size $\leq \frac{n}{2}$, the difference between the anonymous taxation policy and the individual taxation policy is $\left(1 - \frac{s}{n-s}\right)\frac{s}{n}f(n)$.

Now, take some $\alpha \in (0, 1]$; the expression $1 - \frac{s}{n-s}$ is greater than α if and only if $s \leq \frac{1-\alpha}{2-\alpha}n$. For example, setting $\alpha = \frac{1}{2}$ shows that coalitions of size $\leq \frac{n}{3}$ enjoy at least 50% less tax under the optimal taxation policy than under the individual taxation policy, whose cost equals the cost of stability.

We also note that there may be no difference between optimal taxation policies and individual ones, even for anonymous, superadditive games. This is the case, for example, in the three-player majority game; in the three-majority game we have three agents, and the value of all coalitions of size ≥ 2 is 1, while singletons have a value of 0. This game is both anonymous and superadditive, and it has a cost of stability equal to the optimal taxation policy, as it satisfies the conditions set in Theorem 14 presented in the next section.

6. MAXIMAL-STABLE GAMES AND THE LEAST CORE

The previous sections have demonstrated that optimal taxation schemes tend to fare better than individual taxation schemes for rather large classes of games. The purpose of Sections 6 and 7 is to pursue the other direction: suppose that a central authority has implemented a taxation scheme (e.g. the relative least-core); what conditions on the cooperative game must hold in order to ensure that this taxation scheme is indeed minimizing the total tax required in order to stabilize the game? In other words, are there meaningful game classes for which such a taxation scheme does better than the optimal coalitional taxation scheme?

We begin our analysis with finding conditions for maximal-stability in games induced by the various notions of the ε-core described in Section 2.

THEOREM 14. *If $\mathcal{G}_{rel(\varepsilon^*)}$ is maximal-stable then one of the following holds:*

1. *$\varepsilon^* = 1$; i.e. $\mathcal{G}_{rel(\varepsilon^*)} = \mathcal{G}$ and $\mathcal{G}$ has a non-empty core.*
2. *there exists a vector $\mathbf{x} \in \mathbb{R}^n_+$ such that for all $S \subsetneq N$ such that $v(S) > 0$, $v(S) = \frac{1}{\varepsilon^*}\sum_{i\in S} x^i$.*

PROOF. Suppose that $\mathcal{G}_{rel(\varepsilon^*)}$ is maximal-stable; according to Proposition 5, for any $\mathbf{x} \in Core(\mathcal{G}_{rel(\varepsilon^*)})$ and any $S \subseteq N$ we have either $x(S) > \varepsilon^* v(S) = v(S)$, or $x(S) = \varepsilon^* v(S)$. Thus, if there is some $T \subsetneq N$ such that $x(T) > \varepsilon^* v(T)$, and $v(T) > 0$, it must be that $\varepsilon^* = 1$. Otherwise, it must be that for all $S \subsetneq N$ such that $v(S) > 0$ we have that for any outcome $\mathbf{x}$ in $Core(\mathcal{G}_{rel(\varepsilon^*)})$ we have that $x(S) = \varepsilon^* v(S)$. Indeed, this means that for all $S \subset N$ such that $v(S) > 0$, we have that $v(S) = \frac{1}{\varepsilon^*}\sum_{i\in S} x^i$. $\square$

Simply put, Theorem 14 means that the least core taxation scheme is maximal stable only if the game in question is very simple. Similarly, one can show that the standard strong least core induces a maximal-stable game if and only if $Core(\mathcal{G}) \neq \emptyset$ or there exists some $\mathbf{x} \in \mathbb{R}^n$ such that for all $S \subsetneq N$ we have that $v(S) = x(S) + \varepsilon^*$. Finally, for the weak least-core, there exists some $\mathbf{y} \in \mathbb{R}^n$ such that $v(S) = y(S)$ for all $S \subsetneq N$, and the optimal taxation scheme in that case would be to reduce the value of each S by $\varepsilon \cdot |S|$.

While these observations are not too difficult to come by, they do show that for most scenarios of interest, optimal coalitional taxation policies necessarily do better than the common taxation policies existing in the literature. Therefore, there is a significant gap that can be closed via the design of more nuanced taxation schemes.

7. MAXIMAL STABILITY AND THE RELIABILITY EXTENSIONS

In this section, we look at maximal stabiliy vs. reliability extension—a method of incorporating uncertainty into cooperative games. Cooperative reliability games have been explored by Bachrach et al. [3] and are similar to multilinear extensions of cooperative games, introduced by Owen [13]. Similarly to Section 6, we focus on the following problem: given a cooperative game $\mathcal{G}$, what conditions on $\mathcal{G}$ must hold so that there exists some $\mathbf{r} \in [0,1]^n$ such that the reliability extension $\mathcal{G}_\mathbf{r}$ is maximal stable? Theorem 15 below suggests that games for which there exists a maximal stable reliability extension are relatively simple. Namely, these are games for which any coalition that has a positive value with some probability (i.e. all its members have some chance of survival) was additive in the original game.

THEOREM 15. *If there is some $\mathbf{r} \in [0,1]^n$ such that $\mathcal{G}_\mathbf{r}$ is maximal stable, then there exists some $\mathbf{y} \in \mathbb{R}^n$ such that for every $S \subseteq N$ for which $v_\mathbf{r}(S) < v(S)$, and for all $i \in S$ $r^i > 0$, we have that $v(S) = \sum_{i\in S} y^i$.*

PROOF. First note that by Proposition 5, if $v_\mathbf{r}(S) < v(S)$ then there is some $\mathbf{x} \in \mathbb{R}^n$ such that $v_\mathbf{r}(S) = x(S)$. We show that the claim holds by induction on the size of S. If $|S| = 1$ then S is a singleton, so $v_\mathbf{r}(S) = r^i v(\{i\}) = x^i$, and since $r^i > 0$ it must be that $v(\{i\}) = \frac{x^i}{r^i}$. Now, given a set S of size ≥ 2, let us assume that the claim holds for all sets S' such that $|S'| < |S|$. Let us take an arbitrary $i \in S$; we can rewrite $v_\mathbf{r}(S)$ as

$$r^i \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k) v(C\cup\{i\}) + (1-r^i) \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k) v(C)$$

which equals

$$r^i \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k)(v(C\cup\{i\}) - v(C)) + \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k) v(C).$$

Note that the latter sum is simply $v_\mathbf{r}(S\setminus\{i\})$, thus $v_\mathbf{r}(S)$ equals $r^i \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k)(v(C\cup\{i\}) - v(C))$ plus $v_\mathbf{r}(S\setminus\{i\})$, which is equal to $\sum_{j\in S\setminus\{i\}} \frac{x^i}{r^i}$ by induction hypothesis. Also, by induction hypothesis, we have that for all $C \neq S\setminus\{i\}$, it must be that $v(C\cup\{i\}) - v(C) = \frac{x^i}{r^i}$. We add and subtract $\prod_{j\in S} r^j \frac{x^i}{r^i}$ from the sum, to obtain

$$\prod_{j\in S} r^j (v(S) - \sum_{j\in S} \frac{x^j}{r^j}) + r^i \sum_{C\subseteq S\setminus\{i\}} \prod_{j\in C} r^j \prod_{k\in S\setminus\{i\}\setminus C} (1-r^k)\frac{x^i}{r^i}.$$

Observe that for any vector $\mathbf{r} \in [0,1]^n$ and a subset $S \subseteq N$ we have $\sum_{C\subseteq S} \prod_{i\in C} r^i \prod_{j\in S\setminus C} (1-r^j) = 1$. Thus, the right-hand summation in the expression above is equal to $r^i \frac{x^i}{r^i} = x^i$, which in addition to $v_\mathbf{r}(S\setminus\{i\})$ equals $\sum_{i\in S} x^i$. To conclude, we have that $v_\mathbf{r}(S)$ equals

$$x(S) + \prod_{j\in S} r^j \left(v(S) - \sum_{j\in S} \frac{x^j}{r^j} \right).$$

By assumption, $v_\mathbf{r}(S) = x(S)$, which immediately implies that the right-hand expression equals 0; since for all $j \in S$ we have that $r^j > 0$, it must be the case that $v(S) - \sum_{j\in S} \frac{x^j}{r^j} = 0$, which completes the proof. $\square$

We note that if there is some $i \in S$ such that $r^i = 0$, then $v_\mathbf{r}(S) = v_\mathbf{r}(S\setminus\{i\})$; thus, if a reliability extension of a monotone game has $r^i = 0$ for some $i \in N$, then for any coalition S containing i, it must either be the case that $v(S) = v(S\setminus\{i\})$ or that $v(S)$ is some very high value, preventing the stabilization of the game unless it is reduced to 0. To conclude, the types of cooperative games for which reliability extensions induce optimal taxation schemes are not as trivial as those induced by the various notions of ε-core; however, they are mostly additive, with non-additivity implying that the coalition contains an agent whose presence in a coalition guarantees significantly higher revenue. This means that this agent must be consistently excluded from every coalition by setting his survival probability to 0.

8. CONCLUSIONS AND FUTURE WORK

Optimal taxation schemes are highly desirable when attempting to stabilize a cooperative game. In this work, we have shown some properties of games inducing optimal taxation schemes, as well as their relation to other concepts such as the cost of stability, the least core and reliability extensions. Specifically, we obsereved that in general, optimal taxation schemes fare better than minimal individual taxation schemes, where taxes are taken in a non-combinatorial manner. As the latter has a strong connection to the cost of stability (namely, the cost of stability equals the total tax collected under a minimal individual taxation scheme), we explore the connection between optimal taxation schemes and the cost of stability via LP formulations; moreover, in settings where the underlying game is superadditive and anonymous, we show that significatly lower taxes (in the worst case) can be used when implementing an optimal taxation scheme, and show the connection between the tax savings and the size of coalitions. Finally, we explored the opposite direction: rather than comparing optimal taxation schemes to individual taxation schemes, we ascertained the types of games for which a given taxation scheme is optimal. We have found that taxation schemes induced by variants of the ε-core are optimal only for rather trivial games; however, a more complicated class of games admits optimal taxes induced by reliability extensions.

Our understanding of optimal taxation schemes is far from complete. First, the computational complexity of finding an optimal taxation scheme does not seem to fall into the same category as that of other problems such as computing the cost of stability or finding a core element. For example, while deciding whether a given imputation is in the core of a game is co-NP complete, this is not the case for maximal stable games; indeed, in order to decide whether there exists a taxation policy such that a given imputation is in the core of the game after the taxation, one would need to guess a taxation policy, which is an object with no succinct representation in general. Identifying classes of games for which computing optimal taxation policies is in P would be an important step in our understanding of these problems. However, it is not clear how to do so even if one knows that the game admits an optimal taxation scheme induced by, e.g., the reliability extension.

Another issue is the maximal difference between an optimal taxation policy and an individual taxation policy as used to compute the cost of stability. We have completely characterized games for which least core payments coincide with optimal taxation policies, and have shown the relation between the former and the cost of stability. However, it is not immediately clear what is the worst-case loss incurred by utilizing an individual taxation as opposed to taxing whole coalitions. Indeed, we have given partial answers to this question in our work (namely in Sections 6 and 7), but this analysis needs completion

Acknowledgments.

This work was supported by the UK Research Council, under grant EP/1011587/1 for project ORCHID. Yair Zick is additionally supported by SINGA (A*STAR) graduate fellowship.

9. REFERENCES

[1] R. Aumann and J. H. Drèze. Cooperative games with coalition structures. *International Journal of Game Theory*, 3:217–237, 1974.

[2] Y. Bachrach, E. Elkind, R. Meir, D. Pasechnik, M. Zuckerman, J. Rothe, and J. S. Rosenschein. The cost of stability in coalitional games. In *SAGT'09*, pages 122–134, 2009.

[3] Y. Bachrach, R. Meir, M. Feldman, and M. Tennenholtz. Solving cooperative reliability games. In *UAI'11*, pages 27–34, 2011.

[4] C. Bejan and J. Gómez. Core extensions for non-balanced TU-games. *International Journal of Game Theory*, 38:3–16, 2009.

[5] G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan and Claypool, 2011.

[6] M. Davis and M. Maschler. Existence of stable payoff configurations for cooperative games. *Bulletin of the American Mathematical Society*, 69:106–108, 1963.

[7] D. Gillies. *Some Theorems on n-Person Games*. PhD thesis, Princeton University, 1953.

[8] S. Gonzalez and M. Grabisch. Preserving coalitional rationality for non-balanced games. Technical Report halshs-00718358, HAL, 2012.

[9] 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.

[10] R. Meir, J. S. Rosenschein, and E. Malizia. Subsidies, stability, and restricted cooperation in coalitional games. In *IJCAI'11*, pages 301–306, 2011.

[11] R. Meir, Y. Zick, E. Elkind, and J. Rosenschein. Bounding the cost of stability in games with restricted interaction. In *COMSOC'12*, 2012.

[12] R. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, 2(3):225–229, 1977.

[13] G. Owen. Multilinear extensions of games. *Management Science*, 18(5):64–79, 1972.

[14] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*, volume 34 of *Theory and Decision Library. Series C: Game Theory, Mathematical Programming and Operations Research*. Springer, Berlin, second edition, 2007.

[15] D. Schmeidler. The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17:1163–1170, 1969.

[16] A. Schrijver. *Theory of Linear and Integer Programming*. Wiley, Chichester, 1986.