

A Parametric Worst-Case Approach to Fairness in TU-Cooperative Games.

Cosmin BONCHIȘ and Gabriel ISTRATE *

PRELIMINARY AND INCOMPLETE.
PLEASE DO NOT CITE.

Abstract

We study worst-case fairness of a TU-cooperative game, the stable imputation that is most dissimilar to a normative standard of fairness. Motivated by welfare economics, similarity is quantified using information-theoretic divergences. Worst-case fairness aims to parallel the spirit of the price of anarchy from noncooperative game theory in a cooperative setting, quantifying how much deviation from fairness is compatible with coalitional rationality.

Computing our measure is tractable in weighted voting games and many classes of coalitional skill games, but NP-hard in induced-subgraph games and a class of task-count coalitional skill games. In these latter cases we investigate the performance of several approximation algorithms, showing that they yield *constant* approximately optimal solutions. We also upper bound the performance of a Reverse Greedy algorithm on general convex games in terms of two game-specific constants.

JEL classification: C44, C71, D63

Keywords: cooperative games, worst-case fairness, approximation algorithms.

*West University of Timișoara and the eAustria Research Institute, Bd. V. Pârvan 4, cam 045B, Timișoara, RO-300223, Romania. Corresponding author email: gabrielistrate@acm.org

1 Introduction

Stability and *fairness* are two central issues in cooperative game theory. Many of the various solutions proposed emphasize one or the other aspect: the *core* is a prototypical example of a stability-oriented approach, while the Shapley value [33] is the classical, well-understood approach to fairness.

Of these two requirements stability is usually the much less demanding requirement: As long as one can guarantee (perhaps as a result of exogenous interventions such as *subsidies* [6] or *taxation* [41]) that stable imputations exist, to guarantee that such a solution will eventually be chosen one only needs to assume some form of individual or coalitional rationality. In contrast, imposing a normative solution (such as the Shapley value) requires a number of presuppositions:

(a). that the given solution is easy to compute (a statement that is not always true [21]).

(b). that it is the determinate outcome of a (centralized) noncooperative negotiation mechanism, an assumption that is often problematic in a distributed multiagent setting lacking such a regulating mechanism.

(c). that it was not subject to strategic behavior such as *manipulation* [5, 42], nor was it affected by systemic issues such as *agent failures* [8, 7].

(d). that the coalition formation process did not preclude itself (as it may be in the case of successively enlarging coalitions) the adoption of such a solution.

(e). and, finally, that no prior social norms exist in the agent population that favor alternative outcomes [27].

The above objections are, of course, not specific to the Shapley value, and the aim of the previous discussion was suggesting a replacement by any alternative concept: it simply may be the case that no normative approach is appropriate in all circumstances. Assuming stability requires us, however, to confront the issue of solution multiplicity.

In noncooperative game theory the seminal work on the price of anarchy (PoA) [29, 35, 34, 36] provides a powerful alternative to equilibrium selection. Instead of advocating any particular refinement of Nash equilibrium, the price of anarchy measure takes a pessimistic perspective, quantifying the degradation in overall performance due to uncoordinated behavior, measured on the *worst* equilibrium. This circumvents the problem of equilibrium selection by providing (pessimistic) guarantees valid for *any* rational solution.

In this paper we propose an approach with a similar philosophy for TU-

cooperative games. It is fairness in allocations, rather than total coalition payoff that is solution-dependent in this setting. Rather than attempting to postulate any particular “fair cost division”, we investigate the departure from fairness of an arbitrary “rational” cost allocation (where in this paper we define rationality as membership in the core). We measure departure from fairness by employing a parametric family of measures based on variations on the concept of *Rényi entropy*, fruitfully used before as measures of inequality in welfare economics [19].

The structure of the paper is as follows: in Section 2 we discuss related work. In Section 3 we overview some relevant concepts in cooperative game theory, information theory and combinatorial optimization. We then introduce in Section 4 our parametric family of measures of fairness. In Section 5 we discuss the computational complexity of computing worst-case-fairness, highlighting both tractable and NP-hard cases. In Section 6 we present a general approach for obtaining approximately worst-case fair solutions in a convex game, via a “reverse greedy” algorithm whose performance depends on two game-specific constants. We then particularize our discussion to the class of *induced subgraph games* of Deng and Papadimitriou [21].

2 Related work

We are, of course, inspired by the significant amount of work on the price of anarchy [29, 35, 34, 36, 17]. This line of work has inspired a number of related indices that quantify various aspects of noncooperative games. We cite just one such example, the work of Anshelevich et al. [2] on the *price of stability*, that considers the best, rather than the worst Nash equilibrium.

Still in a noncooperative setting, some concepts in the literature address issues related to cooperation, bringing them closer to the scope of the present work. Examples include *coalition-proof Nash equilibria* [10],[11] or the *price of strong anarchy* [1], which restricts the analysis of system behavior to Nash equilibria resilient to deviations by coalitions, as well as the *price of collusion* [26], which measures the inefficiency of the worst possible partition of the set of players. Perhaps the most relevant for this work is the notion of *Price of Democracy* introduced by Chalkiadakis et al. [15]. They attempt (just as we do) to provide a PoA-like measure for cooperative games. However, their setting is different: they explicitly model payoff allocation as a bargaining process (game) and consider the loss in performance in the coalitions arising

from subgame perfect equilibria of this game.

Finally, the study of *non-cooperative justifications of cooperative game theory concepts* is obviously related. We refer the reader to Section 7.1.2 of [16] for a brief outlook.

In contrast to all these results in non-cooperative game theory, our framework makes minimal assumptions with respect to coalition formation and payoff division. Cooperation is not an issue in the examples we discuss: all players are interested in joining the grand coalition (though our framework could be extended to the case of multiple coalitions as well), and no inefficiency arises. What may arise, though, is *inequality* in payoff division.

3 Preliminaries and Notation

All logarithms considered in this paper are base two. We will work in the framework of Cooperative Game Theory (for a recent survey from an algorithmic perspective see [16]). We assume knowledge of basic concepts from this literature. We also assume basic knowledge of computational complexity [4] and approximation algorithms [40].

A *TU-cooperative game* is a pair $\Gamma = (N, v)$, where N is a set of players (usually $N = [n] := \{1, 2, \dots, n\}$ for some $n \geq 1$) and $v : \mathcal{P}(N) \rightarrow \mathbf{R}_+$ is a *value function*. We will assume that v is monotone nonnegative, i.e. $v(\emptyset) = 0$, $[A \subseteq B] \Rightarrow [v(A) \leq v(B)]$. Game Γ is *convex* if the value function v is *supermodular*, that is it satisfies $v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$ for all $A, B \subseteq N$. When the sign of the inequality is reversed function v is called *submodular* and the corresponding game is called *concave*.

An *imputation* is the value received by a coalition of players, if the payoff vector is $x_S = (x_{S_1}, x_{S_2}, \dots, x_{S_k})$, then $x(S) = \sum_{i \in S} x_i$ such that $x(N) = v(N)$. Imputation x is *blocked by coalition* $S \subseteq N$ if $\sum_{i \in S} x_i < v(S)$. The *core of game* Γ is the set of all imputations that are not blocked by any coalition. A *solution concept* q [3] is a function that assigns to every cooperative game Γ an imputation $q(\Gamma)$. It is a *core concept* if $q(\Gamma) \in \text{core}(\Gamma)$ for every game Γ such that $\text{core}(\Gamma) \neq \emptyset$. Basic examples of solution concepts will be U , the uniform vector $U(i) = \frac{v(N)}{|N|}$ for all $i \in N$, as well as *the Shapley value* Sh and *the nucleolus* Nu [16]. In concave games the core is the convex hull of the *marginal vectors* x^π , $\pi \in S_n$ and the Shapley value is the barycenter of this polyhedron.

We compute our measures on three classes of games:

(a) *weighted voting games* (WVG[24]). A WVG is specified by a set of $n \geq 1$ players, a corresponding set of nonnegative player *weights* w_1, w_2, \dots, w_n adding up to 1, and a threshold T . Coalition $S \subseteq [n]$ is *winning* ($v(S) = 1$) if $\sum_{i \in S} w_i \geq T$ and *losing* ($v(S) = 0$) otherwise. Player i is a *veto player* if all winning coalitions must include i .

(b). *coalitional skill games* (CSG[9]). A CSG is specified by a set of *agents*, $N = \{1, \dots, n\}$, a set of *skills* $S = \{s_1, \dots, s_k\}$, a set of *tasks* $T = \{t_1, \dots, t_m\}$ and a *task value function* $u : \mathcal{P}(T) \rightarrow \mathbf{R}_+$. We assume that u is monotone nonnegative, i.e. $u(\emptyset) = 0$, $[T_1 \subseteq T_2] \Rightarrow [u(T_1) \leq u(T_2)]$.

Each agent i has a set of skills $S_i \subseteq S$. Each task t_j requires a set of skills $S_{t_j} \subseteq S$. For $C \subseteq N$, $S(C) = \cup_{j \in C} S_j$ is the *set of skills of coalition* C . Coalition C can perform task t_j if $S_{t_j} \subseteq S(C)$. The set of tasks coalition C can perform will be denoted by $T(C)$. The value function $v(C)$ is defined as $v(C) = u(T(C))$.

In a *single task CSG* (STSG) $T = \{S\}$ hence $v(C) = 1$ if $S(C) = S$, 0 otherwise. Γ is a *task-count CSG* (TCSG) if $u(T') = |T'|$ and a *task-count CSG with threshold* (TCSG-T) if there exists a threshold k such that $u(T') = 1$ if $|T'| \geq k$, 0 otherwise. In a *weighted task-count CSG with threshold* (TCSG-T) the tasks are weighted by some system of nonnegative weights w_1, \dots, w_m and the definition of u is changed to $u(T') = 1$ if $\sum_{l \in T'} w_l \geq k$, 0 otherwise.

(c). *induced subgraph games* (IS-G [21, 38, 16]). An IS-game is specified by a connected loopless graph $G = (V, E)$ and a set of integer *weights* $(w_{i,j})_{(i,j) \in E}$ on the edges. Vertices of G are interpreted as *players*. Given set $S \subseteq V$, the *value of coalition* S is $v(S) = \sum_{(i,j) \in E, i,j \in S} w_{i,j}$. We will assume that weights are nonnegative and for any vertex $v \in V$, the sum of weights of its adjacent edges is positive.

One can normalize every vector of nonnegative values X to a probability distribution. Without risk of confusion we will denote by X the resulting distribution as well. Given discrete random variable X with probability mass function $p = (p_i)_i$ and real number $\lambda > 0$, $\lambda \neq 1$ the *Rényi entropy of order λ of X* is defined [18] as: $H_\lambda(X) = \frac{1}{1-\lambda} \log(\sum_i p_i^\lambda)$. We complete this definition for $\lambda = 1$ by the usual Shannon entropy $H(X) = H_1(X) = -\sum_i p_i \log p_i$. Let $P = (p_i)$ and $Q = (q_i)$ be two distributions and $\lambda > 0$. The *Rényi divergence of order λ of P and Q* is defined as $D_\lambda(P \parallel Q) = \frac{1}{\lambda-1} \log(\sum_i p_i^\lambda q_i^{1-\lambda})$. The *discrete Rényi relative entropy of order λ of P, Q* is defined as $h_\lambda[P, Q] = \frac{1}{1-\lambda} \log(\sum_i q_i^{\lambda-1} p_i) + \frac{1}{\lambda} \log(\sum_i q_i^\lambda) - \frac{1}{\lambda(1-\lambda)} \log(\sum_i p_i^\lambda)$.

It satisfies the *discrete Gibbs inequality*:

Lemma 1. *We have $h_\lambda[P, Q] \geq 0$.*

Lemma 1 is the discrete version of a result from [30]. Its proof will be given in the full version.

We complete the two definitions above in the special case $\lambda = 1$ by the *relative Shannon entropy* or *Kullback-Leibler divergence*, defined as: $D(P \parallel Q) = h_1[P, Q] = \sum_i p_i \log \frac{p_i}{q_i}$.

Though they do not generally yield metrics, entropy and divergence measures have a significant history (e.g. [39]) of use, in particular as indicators of "similarity" or "distance" between two probability distributions. A particularly important application of such information-theoretic tools is in the area of *inequality measurement* [37, 19]. This is what motivates the use of Rényi divergences as objective function to maximize in our measure, that we define next.

4 Worst-case fairness of a TU-cooperative game

We now define the main object of interest, a parametric family of measures of fairness for cost allocations of a TU-cooperative game $\Gamma = (N, v)$. They are parameterized by

- (a). a positive real λ .
- (b). A set $St(\Gamma)$ of solutions deemed "stable". In all examples considered in this paper $St(\Gamma) = Core(\Gamma)$.
- (c). a solution concept q , yielding vector $q(\Gamma) \in \mathbf{R}_+^{|N|}$. Intuitively vector $q(\Gamma)$ represents a baseline "standard of fairness" to which all other possible imputations are held. Our measures attempt to evaluate the largest possible discrepancy between a stable imputation $u \in St(\Gamma)$ and $q(\Gamma)$.

These considerations finally enable us to give the definition of worst-case fairness: Given cooperative game $\Gamma = (N, v)$ and real number $\lambda > 0$ the λ -*worst-case fairness of game* Γ with respect to (St, q) is defined as

$$OPT_\lambda(\Gamma, St, q) = \sup\{D_\lambda(x||q(\Gamma)) : x \in St(\Gamma)\}. \quad (1)$$

Generally we will not be content with only computing the optimal value in equation (1), but instead also seek to compute (if possible) a vector $W(\Gamma)$

that realizes equality

$D_\lambda(W(\Gamma)||q(\Gamma)) = OPT_\lambda(\Gamma, St, q)$. Such a $W(\Gamma)$ will be called a λ -worst-case fair imputation of Γ .

The definition in equation (1) obviously depends on the choice of q . Several special cases make sense:

(a). *strictly egalitarian worst-case fairness*: $q(\Gamma)$ is the uniform vector U . Though somewhat controversial since it requires a very strong form of equality, the study of this measure makes sense at least from a *mechanism design* point of view: To give just one example, in the case of convex/concave games, particularly interesting examples of imputations in the core arise from group-strategyproof mechanisms or, equivalently, cross-monotonic sharing schemes (see [31] and Chapter 15 of [32]). Requiring cross-monotonicity yields a “plausible notion of equity” [28]. A natural question related to the previous quote is how large a variation in payoffs is compatible with the use of cross-monotonic schemes. The strictly egalitarian WCF offers a pessimistic estimate of this amount.

(b). *marginalist worst-case fairness*: in this case q is (the probability distribution obtained from) the Shapley value.

(c). *lexicographic worst-case fairness*: q is (the probability distribution obtained from) the nucleolus of the game.

(d). *egalitarian worst-case fairness*: in this case q corresponds to the egalitarian solution of Dutta and Ray [23]. We will not study this measure in the present paper.

Worst case fairness is obviously easy to compute when (i). The baseline concept $q(\Gamma)$ is tractable. (ii). a witness $W(G)$ is easily computable. We will limit ourselves in this paper to instances where condition (i). holds, so that the (in)tractability of computing the worst-case fairness is not due to the intractability of the baseline concept.

Example 1. Consider the IS-game Γ with three players presented in Figure 1(a). The total payoff to be shared between players is $12 = 2 + 4 + 6$. The core of Γ is given in Figure 1(b). The Shapley value is [21] $Sh = (3; 4; 5)$. Furthermore, by simple symbolic computations one can show the following: (a). There are 57 integral imputations in the core. Of these six are extremal (i.e. corners of the polyhedron geometrically describing the core). (b). For $\lambda = 1$ there are two WCF imputations in the core for the strictly egalitarian WCF: $(2; 0; 10)$ and $(0; 2; 10)$. Worst-case fairness is $OPT_1(\Gamma, U) \sim 0.934$. (c). Also for $\lambda = 1$ there exists an unique WCF imputation for the marginal-




Figure 1: (a) An IS game with three players; (b) Its core.

ist WCF : $X = (4; 8; 0)$. The corresponding WCF is $OPT_1(\Gamma, Sh) \sim 0.805$. Thus employing perfect equality as our standard of fairness of imputations in the core is (in this game) more pessimistic than using the Shapley value.

5 The computational complexity of worst-case fairness

In this section we study the computational complexity of computing the WCF of a TU-game. For technical reasons (Rényi entropy is only concave for $\lambda \leq 1$) our intractability results will be stated under the condition $0 < \lambda \leq 1$. On the other hand all positive results (including those related to approximate solutions) deal with the general case $\lambda > 0$.

First we provide two settings where computing the WCF is tractable, that of weighted voting games and the one of coalitional skill games.

Theorem 1. *Let $\Gamma = ([n], \{w_i\}_{i \in [n]})$ be a weighted voting game with nonempty core. Without loss of generality assume that players $1, 2, \dots, K$ are all the veto players (to have a nonempty core necessarily $K \geq 1$).*

Then for every $\lambda > 0$ vector $P = (1, 0, \dots, 0)$ is a WCF-imputation for the (strictly egalitarian/lexicographic) λ -worst case fairness in Γ .

Theorem 2. *Let $\Gamma = (I, S, T, u)$ be a STSG (TCSG-T, WTSG-T) game with nonempty core. Without loss of generality assume that players $1, 2, \dots, K$ are all the veto players (to have a nonempty core necessarily $K \geq 1$). Then for every $\lambda > 0$ vector $P = (1, 0, \dots, 0)$ is a WCF-imputation for the strictly egalitarian λ -worst case fairness in Γ .*

Moreover, if Γ is a TCSG-T game then the same conclusion holds for the lexicographic worst-case fairness of Γ .

Proof. The two theorems above have a common proof, based on the similar characterization of the core of weighted voting games [24] and coalitional skill games [9]: an imputation $x = (x_1, x_2, \dots, x_n)$ is in the core if it distributes zero value to every non-veto player. That is

$$x_1 + x_2 + \dots + x_K = 1 \text{ and } x_{K+1} = \dots = x_N = 0.$$

On the other hand, for WVG and TCSG-T games the nucleolus of Γ is the vector $Nu(\Gamma) = (\frac{1}{K}, \dots, \frac{1}{K}, 0, \dots, 0)$.

As for all vectors x in the core $x_i = 0$ for nonveto players i , and since vectors $U, Nu(\Gamma)$ are constant on their restriction to veto players, maximizing divergences $D_\lambda(x||U)$, and $D_\lambda(x||Nu)$ respectively, is easily seen to be equivalent to minimizing the entropy of the distribution of payoffs restricted to veto players $\bar{x} = (x_1, x_2, \dots, x_k)$. Obviously, vector P is one way to accomplish this task. \square

The case of TCSG-games is rather different: while of the three target concepts U, Sh, Nu only U is clearly tractable (Sh is \sharp P-complete and the complexity of Nu is open [9]), computing even the strictly egalitarian WCF is NP-hard:

Theorem 3. *For any $0 < \lambda \leq 1$ the following decision problems is NP-hard: Given a TCSG game Γ and a constant $\eta > 0$, does there exist any imputation $x \in core(\Gamma)$ with $D_\lambda(x, U) \geq \eta$?*

Proof. Consider the following subclass of TCSG games: a *pure skill CSG game* (PCSG) is a TCSG game for which $T = S$. We will show that the problem above remains NP-hard even for PCSG games.

Indeed, every PCSG game Γ is easily seen to be concave¹. An imputation of Γ may (fractionally) divide each skill among players who possess it. By stochastic domination, to maximize $D_\lambda(X||U)$ (i.e. minimize $Ent_\lambda(X)$ ($= \log(n) - D_\lambda(X||U)$)) one only needs to consider integral divisions (each skill is assigned to a player). This is essentially the problem *Minimum Entropy Set Cover* (MESC [25]). Problem MESC is NP-hard (the general case $0 < \lambda \leq 1$ is dealt with essentially in [12]). \square

Finally, in IS-games the nucleolus is equal to the Shapley value, the latter one being easily computable [21]. Computing worst-case fairness in IS games, both strictly egalitarian and marginalist, is computationally intractable:

¹this is not true for general TCSG games

Theorem 4. *For any $0 < \lambda \leq 1$ the following decision problems are NP-hard:*

- (a). *Given an IS-game $\Gamma = (G, w)$ with nonnegative weights and constant $\eta > 0$, does there exist any imputation $x \in \text{core}(\Gamma)$ with $D_\lambda(x, U) \geq \eta$?*
- (b). *Given an IS-game $\Gamma = (G, w)$ with nonnegative weights and constant $\eta > 0$, does there exist any imputation $x \in \text{core}(\Gamma)$ with $D_\lambda(x, Sh) \geq \eta$?*

Note that we have framed our intractability results above in terms of general *real-valued* constants λ, η . With only additional notational inconvenience we can make all these constants rational: the set of potentially optimal solutions to all instances of the games we have considered above is finite and the objective values "well-spaced".

6 Approximation algorithms

Given Theorem 4, we need to give up (at least in IS-games) the hope of computing efficiently computing WCF-imputations, and instead resort to *approximation algorithms*, that will provide *approximately worst-case fair* imputations. Unlike most cases in the theory of approximation algorithms [40], but similar to other problems in entropy minimization [14], our approximation guarantees will be *additive*: given constant $\Delta > 0$, an imputation X will be Δ -*approximately worst-case fair* if

$$D_\lambda(X||q(\Gamma)) \geq \text{OPT}_\lambda(\Gamma, St, q) - \Delta$$

and we will try to find imputations minimizing Δ .

It is customary when dealing with approximation in submodular optimization to employ the GREEDY algorithm, that enlarges the coalition by selecting the individual with the largest marginal increase of the value function v .

For instance in PCSG-games, taking into account the connection with MESC from the proof of Theorem 3, a relatively straightforward adaptation of the argument in [12] (for $0 < \lambda \leq 1$), or as an application of results in Section 7 to the dual of the game (this argument works for all $\lambda > 0$) yields the following result (proved in the full version of the paper):

Theorem 5. *Let $\lambda > 0$ and Γ be a PCSG-game. Then the GREEDY algorithm produces a $\frac{1}{\lambda-1} \log(\lambda)$ -approximately worst-case-fair imputation for the strictly egalitarian WCF in Γ . The approximation guarantee is optimal unless $P=NP$.*

On the other hand, for games (such as IS-games) that are *supermodular*, using the GREEDY algorithm doesn't quite make sense: assigning the first element i its payoff $v(\{i\})$ does not take into account the fact that the contribution of player i increases with the coalition, being largest for the coalition $N \setminus \{i\}$. In other words $v(\{i\}) \leq v(N) - v(N \setminus \{i\})$, and, to create an imbalanced allocation we should assign player i its utopia payoff (that is the right-hand quantity, rather than the left-hand). This leads to considering the Reverse Greedy algorithm displayed below.

Reverse Greedy:

```

INPUT : A game  $\Gamma = (N, v)$ 
 $y := (0, 0, \dots, 0)$ 
 $A_0 := N, r:=1$ 
While  $\exists e \in A_{r-1}$  with  $v(A_{r-1}) - v(A_{r-1} \setminus \{e\}) > 0$ 
    choose  $i_r \in A_{r-1}$  that maximizes
     $v(A_{r-1}) - v(A_{r-1} \setminus \{i_r\})$ 
    (breaking ties arbitrarily)
     $y_{i_r} := v(A_{r-1}) - v(A_{r-1} \setminus \{i_r\})$ 
     $A_r := A_{r-1} \setminus \{i_r\}, r++$ 

OUTPUT : Imputation  $Y = (y_i)_{i \in N}$ .

```

Figure 2: Algorithm Reverse Greedy.

Example 2. *Consider the setting of Example 1. Algorithm Reverse Greedy computes one of the covers $(0; 2; 10)$ or $(2; 0; 10)$ (optimal for strictly egalitarian WCF). The computed imputation depends on the tie-breaking rule between the first two nodes. Indeed, the algorithm first selects node C , allocating its utopia value $4 + 6 = 10$. Then it selects one of A and B in an arbitrary order.*

For IS-games we will consider an alternate approximation rule: we will call an imputation BI *biased* if for every edge $(i, j) \in E$ it distributes all

the weight $w_{i,j}$ to one of the nodes i, j , when $w(i) \neq w(j)$ to the the node among i, j with larger value of $w(\cdot)$. Clearly, a biased imputation is easy to compute.

Example 3. *In the setting of Examples 1, 2, imputation $(0; 2; 10)$ is the only biased imputation.*

7 Algorithm Reverse Greedy in arbitrary concave games

Our main result yields an upper bound on the performance of Algorithm Reverse Greedy in approximating the strictly egalitarian worst-case fairness in an arbitrary convex TU-game. This easily yields (via Lemma 2 below) a weaker additive guarantee for any solution concept (we will take this route in the next section)

To describe our guarantees we introduce some notation:

1. We will denote by l the number of iterations of the Reverse Greedy algorithm.
2. For $1 \leq r \leq l$ denote by i_r the element chosen at stage r of the algorithm. Let $W_r = \{i_1, \dots, i_r\}$, $A_r = U \setminus W_r$ and Δ_r be the value of element y_{i_r} set at stage r .

We next define a quantity, the "impact of j on i_r ", that will play a fundamental role in our results below: For any $1 \leq r \leq l$ we define the impact of j on i_r by

$$a_r^j = [v(A_{r-1}) - v(A_r)] - [v(A_{r-1} \setminus \{j\}) - v(A_r \setminus \{j\})]. \quad (2)$$

Proposition 1. *For any $1 \leq r \leq l$ and $1 \leq j \leq m$ we have $a_r^j \geq 0$.*

Proof. Note that $A_{r-1} = A_r \cup \{i_r\}$. Thus when $j = i_r$ or $i_r \neq j \notin A_{r-1}$ the second term is zero, and the result follows directly from the monotonicity of function v . Assume now that $i_r \neq j \in A_{r-1}$, thus $j \in A_r$. Define $S = A_r$ and $T = A_{r-1} \setminus \{j\}$. Then $S \cup T = A_{r-1}$, $S \cap T = A_r \setminus \{j\}$, and we employ the supermodularity of function v . \square

Given an optimal solution $X = (X_j)$, we will break it down into a large number of components $Z_r^j \in \mathbf{Z}, 0 \leq Z_r^j \leq a_r^j$ as in equation (3) below:

$$X_j = \sum_{r=1}^l Z_r^j, \forall j \in [m] \quad (3)$$

Intuitively Z_r^j is the part of the optimal solution X_j that can be assigned to cover “set i_r ”. This explains the newly introduced constants: first, one cannot allocate more than the total of X_j . Second, one cannot allocate to any “set i_r ” more than “its intersection with X_j ”.

Definition 1. *Given concave game Γ Let $\alpha = \alpha(\Gamma)$ the smallest and $\beta = \beta(\Gamma)$ the largest positive value such that for some cover X in the core minimizing $Fair_\lambda(\Gamma)$, one can define quantities Z_r^j , so that for any $r \in [l]$:*

$$\beta \cdot \Delta_r \leq \sum_{j=1}^m Z_r^j \leq \alpha \cdot \Delta_r. \quad (4)$$

Proposition 2. *For any convex game Γ we have*

$$\beta(\Gamma) \leq 1 \leq \alpha(\Gamma).$$

Proof. We prove first inequality, the second is similar.

Sum equations (4) for $r = 1, \dots, l$. The left-hand side is $\beta(\Gamma) \sum_{r=1}^l \Delta_r = \beta(\Gamma) f(N)$, by ReverseGreedy.

Similarly, the right-hand side is $\sum_{r=1}^l \left(\sum_{j=1}^m Z_r^j \right) = \sum_{j=1}^m \left(\sum_{r=1}^l Z_r^j \right) = \sum_{j=1}^m X_j = f(N)$. The result follows. $\square \quad \square$

Our main result gives an upper bound applicable to all convex cooperative games:

Theorem 6. *Given a convex cooperative game Γ the Reverse Greedy algorithm produces a cover RG satisfying*

$$0 < \lambda < 1 \Rightarrow H_\lambda(RG) \leq H_\lambda(OPT) + \frac{1}{\lambda - 1} \log(\beta\lambda). \quad (5)$$

$$\lambda > 1 \Rightarrow H_\lambda(RG) \leq H_\lambda(OPT) + \frac{1}{\lambda - 1} \log(\alpha\lambda). \quad (6)$$

Corollary 1. *The cover RG produced by the ReverseGreedy algorithm is $\frac{1}{\lambda-1} \log(\beta\lambda)$ -approximately WCF with respect to U for $0 < \lambda < 1$ and $\frac{1}{\lambda-1} \log(\alpha\lambda)$ -approximately WCF with respect to U for $\lambda > 1$*

Proof. Follows directly from Theorem 6. □

Observation 1. *By Proposition 2 both constants in the upper bounds of Theorem 6 are nonnegative. On the other hand, if at least one of parameters α or β are equal to 1 then we can complete the result to the case $\lambda = 1$ by taking the limit $\lambda \rightarrow 1$, yielding the conclusion that RG is $\log(e)$ -approximately fair with respect to U for $\lambda = 1$.*

A more limited connection between divergence and entropy holds even in the general case. Given distribution $R = (r_i)$, denote $r_{max} = \max\{r_j : j \in \text{supp}(R)\}$, $r_{min} = \min\{r_j : j \in \text{supp}(R)\}$ and define $\nu(R) = \log\left(\frac{r_{max}}{r_{min}}\right)$, the nonuniformity of distribution R .

Lemma 2. *Let P, Q, R be probability distributions and $\lambda > 0$. Then $|D_\lambda(P||R) - D_\lambda(Q||R) - (H_\lambda(Q) - H_\lambda(P))| \leq \nu(R)$.*

The proof of Lemma 2 is deferred to the full version. We will apply it to IS-games below.

8 Approximate Worst-case fairness of Induced Subgraph Games

In this section we first particularize our main result to the class of IS-games: we show that for any such game $\alpha = \beta = 1$. On the other hand we study the performance of using biased imputations, showing that in some cases (for $\lambda \sim 1$) the guarantee is better than the one available for the ReverseGreedy algorithm:

Theorem 7. *Given IS game $\Gamma = (G, w)$ and $\lambda > 0$*

(a). $\alpha(\Gamma) = \beta(\Gamma) = 1$.

(b). *Any biased imputation BI satisfies*

$$H_\lambda(Sh) - H_\lambda(OPT) \leq \frac{1}{\lambda} [H_\lambda(Sh) - H_\lambda(BI)] + 1. \quad (7)$$

Corollary 2. *In the setting of the previous result RG is $\frac{1}{\lambda-1} \log(\lambda)$ -approximately worst-case fair with respect to U and $\frac{1}{\lambda-1} \log(\lambda) + \nu(Sh)$ -worst case fair with respect to Sh .*

On the other hand BI satisfies

$$D_\lambda(BI \parallel Sh) \geq \lambda \cdot OPT_\lambda(\Gamma, Sh) - (1 + \lambda) \cdot \nu(Sh) - \lambda \quad (8)$$

Proof. Directly from Theorem 7 and Lemma 2. □

For marginalist worst-case fairness the second bound may be slightly better when $\lambda \approx 1$ and $\nu(Sh) \approx 0$. Indeed, in the limit $\lambda \rightarrow 1$ term $\frac{1}{\lambda-1} \log(\lambda)$ tends to $\log(e) \approx 1.442\dots$, while $\lambda \approx 1$. The best of the two guarantees may depend on the precise value of constant λ (and, of course, other features of the instance at hand).

9 Delayed Proofs

9.1 Proof sketch of Theorem 4

For reasons of space the proof of this result is only outlined. A complete argument is deferred to the full version.

The first ingredient of our proof is the characterization of imputations $Z = (z_1, z_2, \dots, z_n)$ in the core of an IS game:

Proposition 3. *z is in the core of Γ if and only if there exist real numbers $r_{i,j}$ are real numbers in the range $0 \leq r_{i,j} \leq 1$ with $r_{i,j} + r_{j,i} = 1$ so that*

$$z_i = \sum_{(i,j) \in E} r_{i,j} w_{(i,j)} \quad (9)$$

This claim is an easy consequence of the characterization of the core of IS games [21], and a special case of a more general paradigm [20].

Lemma 3. *Let $X = (x_i)_i$ be an optimal imputation in game Γ . Consider a reordering σ of the set of vertices so that $x_{\sigma(1)} \geq x_{\sigma(2)} \geq \dots \geq x_{\sigma(n)}$. Then coefficients $r_{k,l}$ from formula (9) satisfy: $r_{\sigma(i),\sigma(j)} = 1$, if $i < j$, 0 otherwise, for all $i, j \in V$.*




Figure 3: The gadget in [9] and our reductions.

Proof. Consider an arbitrary edge $(k, l) \in E$, $k = \sigma(i)$, $l = \sigma(j)$. Assume that $i < j$ (the opposite case is easily handled via relation $r_{k,l} + r_{l,k} = 1$) and $r_{k,l} < 1$. Define for notational convenience $0 < \epsilon < 1$ by $\epsilon = r_{l,k} = 1 - r_{k,l}$. With this choice further define $\tilde{X} = (\tilde{x}_1, \dots, \tilde{x}_k, \dots, \tilde{x}_l, \dots, \tilde{x}_n)$ where $\tilde{x}_k = x_k + \epsilon w_{k,l}$, $\tilde{x}_l = x_l - \epsilon w_{k,l}$, and $\tilde{x}_r = x_r$ for all other $r \neq k, l$. Note that \tilde{x} and x differ just on components k, l . By the previous remark \tilde{x} is an imputation in the core.

It is easy to see that for all $\lambda > 0$ we have $H_\lambda(X) > H_\lambda(\tilde{X})$, which contradicts the hypothesis that X had the lowest Rényi entropy. \square

Another way to state Lemma 3 is that any imputation of minimal entropy corresponds to "orienting" the weighted edge (i, j) towards one of the nodes, that is assigning one of the nodes the entire weight $w_{i,j}$.

The proof of the two parts of Theorem 4 are fairly similar, and mirror the NP-hardness proof (given in [13]) of a problem called *minimum entropy orientation* (MINEO).

(a). By Lemma 3 our problem is equivalent (in the sense of having the same optimum, though the set of feasible solutions may vary) to a weighted version of MINEO.

The reduction in [13] encodes an instance Φ of an NP-complete variant of problem 1-in-3 SAT into a graph $G = (V, E)$. Formula Φ (corresponding to a set cover problem) has ([13]) $q \geq 1$ variables u_1, u_2, \dots, u_q and an equal number of clauses. Each variable occurs in exactly 3 clauses of Φ , each clause has length exactly 3. Graph G is constructed as the union of q "gadgets" (each having six vertices as displayed in Figure 3 - see [13]) and q extra nodes. It has $m = 12q$ edges.

We may read the existence of a satisfying assignment for Φ from the value of a minimum entropy orientation \vec{H} :

Lemma 4. *The Shannon entropy of the distribution corresponding to an arbitrary orientation \vec{H} of G is at least $\frac{1}{m}(4q \log(m/4) + 7q \log(m/3) + q \log(m))$ with equality (reached for some orientation \vec{G} explicitly constructed in [13]) if and only if instance Φ is satisfiable.*

The proof of Theorem 4 (a) is a direct extension to the case $\lambda > 0$ of the above result. We replace the Shannon entropy by the Rényi entropy, i.e. $\log(n)$ minus the quantity we attempt to maximize, the divergence with the uniform distribution U . Instead of Claim 2 in [13] we use the following extension to all values $\lambda > 0$, with essentially the same proof as the original version:

Lemma 5. *For any $\lambda > 0$, $\lambda \neq 1$, the Rényi entropy of the distribution corresponding to orientation \vec{G} (constructed in [13]) is at most $\log(m) + \frac{1}{1-\lambda} \log(4^\lambda + 7 \cdot 3^{\lambda-1} + 1) / 12$, with equality if and only if instance Φ is satisfiable.*

(b). For simplicity we outline here the case $\lambda = 1$. As in (a) there is no problem extending it to a general $\lambda > 0$.

We employ the same reduction from [13] of 1-in-3 SAT to an instance of the decision problem in Theorem 4 (b): given instance Φ we construct graph G and explicit constant μ_0 such that for every orientation in G $D_\lambda(H, Sh) \leq \mu_0$. Equality can be reached **if and only** if formula is satisfiable.

What changes is the proof of the correctness of the reduction. The nature of the Shapley value forces this: Constructed graph G has two types of nodes, those having degree three and those having degree four. Deng and Papadimitriou [21] proved that the Shapley value of an IS-game is $s(i) = \frac{1}{2} \sum_{i \neq j} w_{i,j}$. Thus, in the unweighted IS-game on G the Shapley node i has Shapley value $s(i) = \frac{3}{2}$ if its degree is three and $s(i) = 2$ if the degree is four.

The following simple lemma shows how to orient edges in an arbitrary orientation in order to maximize divergence with the distribution $Sh = (s_i)$, where $s_i = s(i)/m$.

Lemma 6. *Consider an orientation H in graph G and let A, D be (Figure 4) two connected nodes, one having degree four, the other degree 3. Let a, d be*




Figure 4: An edge between vertices with different degrees.

the indegrees of nodes A, D **not including edge AD** . Then if $a > d$, to maximize the Kullback-Leibler divergence of H with Sh , edge AD should be oriented towards A . Otherwise it should be oriented towards D .

The next step of the construction in [13] (Claim 1 in that paper) was to determine the shape of a minimum entropy (maximum divergence) orientation in graph G on a six-vertex gadget X . Our proof employs a similar lemma. However, since in any gadget X the outside node has degree four and the "interior nodes" have degree three, we need to be more precise, and identify in the maximum divergence configuration the degrees of the interior nodes in X and those of the exterior nodes. We defer a precise statement/proof of our analog of Claim 1 to the final version. The proof is not difficult (it uses Lemma 6) but rather long and computational/cumbersome.

An outcome of the analysis of is the following: in any orientation H maximizing $D(H, Sh)$ the outdegree of nodes A, B, C in every copy X of the gadget is either one or two. Let t be the number of gadgets with indegree 1 (the rest of them, $q - t$ having indegree two). The corresponding local configurations of maximal divergence have the form $O = (4, 1, 0), I = (0, 3, 2)$, if $\text{indeg}(X) = 1$, and $O = (4, 3, 0), I = (0, 1, 3)$, if $\text{indeg}(X) = 2$.

Finally, Theorem 4 (b). follows from the following

Claim 1. *The Kullback-Leibler divergence $KL(\vec{G} || Sh)$ between any orientation \vec{G} and Sh , the probability distribution corresponding to the Shapley value is at most $\frac{q}{m} (9 + 2 \log \frac{3}{2})$, with equality if and only if instance Φ is satisfiable.*

To prove the claim and the theorem we consider the orientation H maximizing $D(H, Sh)$ and note that the contribution of non-gadget nodes to $D(H, Sh)$ is entropy-like (Sh is constant on non-gadget nodes, as all of them have degree three). The sum of degrees of these nodes is $q + t (= (q - t) + 2 \cdot t)$. By stochastic domination we maximize $D(H, Sh)$ by making as many degree three nodes as possible (i.e. $\lfloor \frac{q+t}{3} \rfloor$ nodes). Finally, to obtain the orientation

of maximal divergence we maximize the "optimal divergence" over t . This happens for $t = 0$. That is, just as in the corresponding result in [13], every copy of the gadget has outdegree 1 and the outgoing edges have to give in-degree three to corresponding non-gadget nodes. This is only possible when formula Φ is satisfiable .

9.2 Proof of Theorem 6

Denote by $OPT = (X_j)_{j \in [n]}$ and $RG = (y_i)_{i \in [n]}$ the optimal solution, respectively the one generated by the Reverse Greedy algorithm.

For $1 \leq r \leq l$ we will use the shorthand $U_j^r = X_j - \sum_{k=1}^r Z_k^j$ and $U_j^0 = X_j$. For any fixed j , sequence (U_j^r) is decreasing with r . On the other hand $U_j^{r-1} - U_j^r = Z_r^j$.

By the greedy choice we infer $y_{i_r} = \Delta_r$ with $y_i = 0$ for other values of i . Starting from $A_0 = N$ we have:

$$\begin{aligned} \Delta_r &\geq f(A_{r-1}) - f(A_{r-1} \setminus \{j\}) = f(N) - [f(N) - f(A_{r-1})] + \\ &\quad + [f(N \setminus \{j\}) - f(A_{r-1} \setminus \{j\})] - f(N \setminus \{j\}) = f(N) - \\ &\quad - \sum_{k=1}^{r-1} [f(A_{k-1}) - f(A_k)] + \sum_{k=1}^{r-1} [f(A_{k-1} \setminus \{j\}) - f(A_k \setminus \{j\})] - \\ &\quad - f(N \setminus \{j\}) \geq f(N) - f(N \setminus \{j\}) - \sum_{k=1}^{r-1} a_k^j \geq X_j - \sum_{k=1}^{r-1} a_k^j. \end{aligned}$$

At the last step we used inequality $X_j \leq f(N) - f(N \setminus \{j\})$, which follows from core membership (in)equalities

$$\sum_{k \in N \setminus \{j\}} X_k \geq f(N \setminus \{j\}) \text{ and } \sum_{k \in N} X_k = f(N).$$

Case $\lambda > 1$:

First we use inequality $\sum_{j=1}^m Z_r^j \leq \alpha \cdot \Delta_r$ as follows:

$$\alpha \sum_{r=1}^l (\Delta_r)^\lambda = \sum_{r=1}^l (\alpha \Delta_r) (\Delta_r)^{\lambda-1} \geq \sum_{r=1}^l \left(\sum_{j=1}^m Z_r^j \right) \Delta_r^{\lambda-1}$$

Applying the lower bound above on Δ_r we get:

$$\begin{aligned}
\sum_{r=1}^l \sum_{j=1}^m Z_r^j \Delta_r^{\lambda-1} &\geq \sum_{r=1}^l \sum_{j=1}^m Z_r^j \left(X_j - \sum_{k=1}^{r-1} Z_k^j \right)^{\lambda-1} = \\
\sum_{j=1}^m \left[\sum_{r=1}^l Z_r^j (U_j^{r-1})^{\lambda-1} \right] &= \sum_{j=1}^m \sum_{r=1}^l Z_r^j (U_j^{r-1})^{\lambda-1} = \\
&= \sum_{j=1}^m \sum_{r=1}^l (U_j^r - U_j^{r-1}) (U_j^{r-1})^{\lambda-1} (*)
\end{aligned}$$

We transform the difference into a sum of ones. As $x^{\lambda-1}$ is increasing and $U_j^0 = X_j$ we can lower bound (*) by:

$$\sum_{j=1}^m \sum_{r=1}^l \sum_{k=U_j^r+1}^{U_j^{r-1}} (U_j^{r-1})^{\lambda-1} \geq \sum_{j=1}^m \sum_{r=1}^l \sum_{k=U_j^r+1}^{U_j^{r-1}} k^{\lambda-1} = \sum_{j=1}^m \sum_{k=1}^{X_j} k^{\lambda-1}$$

Putting things together, using standard calculus:

$$\alpha \sum_{r=1}^l (\Delta_r)^\lambda \geq \sum_{j=1}^m \left(\sum_{k=1}^{X_j} k^{\lambda-1} \right) \geq \sum_{j=1}^m \frac{X_j^\lambda}{\lambda} = \frac{1}{\lambda} \sum_{j=1}^m X_j^\lambda$$

Taking the logarithm and dividing by $1 - \lambda < 0$ yields:

$$\frac{1}{1-\lambda} \log \left(\sum_{r=1}^l \Delta_r^\lambda \right) \leq \frac{1}{1-\lambda} \log \left(\sum_{j=1}^m X_j^\lambda \right) - \frac{1}{1-\lambda} \log(\alpha\lambda)$$

or, equivalently, by the definition of Rényi entropy:

$$H_\lambda(RG) \leq H_\lambda(OPT) + \frac{1}{\lambda-1} \log(\alpha\lambda)$$

The proof is similar in the case $0 < \lambda < 1$. We use instead the definition of β . Also the standard calculus inequality changes its direction.

9.3 Proof of Theorem 7 (a)

We first reinterpret the result of Lemma 3 as follows: any optimal solution $X = (x_i)_i$ corresponds to some ordering σ of the vertices such that

$$x_i = \sum_{\substack{(i,j) \in E \\ \sigma^{-1}(i) < \sigma^{-1}(j)}} w_{(i,j)}.$$

Lemma 7. *Given any IS game (G, w) we have*

$$a_r^j = \begin{cases} w_{i_r, j}, & \text{if } i_r \neq j, (i_r, j) \in E, j \in A_r \\ \Delta_r, & \text{if } i_r = j \\ 0, & \text{otherwise,} \end{cases} \quad (10)$$

where Δ_r is the value computed by the algorithm *Reverse Greedy* at stage r .

Proof. A simple application of formulas defining coefficient a_r^j : in IS-games for any set $S \subseteq V$, $v(S) = \sum_{e \in S \times S} w_e$. Therefore $v(A_{r-1}) - v(A_r)$ is the sum of weights w_e of edges e between i_r and a node in A_r . On the other hand the value of expression $v(A_{r-1} \setminus \{j\}) - v(A_r \setminus \{j\})$ depends on j : It is zero if $i_r = j$, $v(A_{r-1}) - v(A_r)$ when $j \neq i_r$ and $j \notin A_r$, otherwise, it is equal to the sum of weights w_e of edges e between i_r and a node in $A_r \setminus \{j\}$. In particular this is $v(A_{r-1}) - v(A_r)$ when j is not adjacent to i_r . \square

Lemma 8. *For all IS games $\Gamma = (G, w)$ one can construct a system of parameters (Z_r^j) from Equation (3), witnessing equality $\alpha(\Gamma) = \beta(\Gamma) = 1$.*

Proof. Lemma 7 allows us to define system of coefficients (Z_r^j) s.t. for all r , $\sum_j Z_r^j = \Delta_r$. Together with Definition (1) and Proposition (2) this witnesses the fact that $\alpha(\Gamma) = \beta(\Gamma) = 1$.

In the construction we will regard OPT and RG as edge orientations in the weighted graph (G, w) . Note that Δ_r is the sum of weights of all edges oriented towards i_r in RG. Intuitively we redistribute this amount among coefficients Z_r^j with $1 \leq j \leq m$ by comparing orientations OPT and RG. Edges are considered in the order given by *Reverse Greedy*:

- Start with $Z_r^j = 0$ for all r and j .
- Run algorithm *Reverse Greedy* that constructs orientation RG, updating coefficients during the algorithm:

- **At each stage r :** after choice of vertex i_r we consider edges (i_r, j) oriented by Reverse Greedy towards i_r . There are two possibilities:
 1. (i_r, j) is oriented towards i_r in both RG and OPT. We set $Z_r^{i_r} = Z_r^{i_r} + w_{i_r, j}$.
 2. (i_r, j) is oriented differently in OPT and RG. We let $Z_r^j = w_{i_r, j} (= a_r^j)$ for such edges.

Note that the total weight assigned at stage r is $\Delta_r (= a_r^{i_r}$ according to Lemma 7).

Hence inequality $0 \leq Z_r^j \leq a_r^j$ is true for $j = i_r$ too.

This completes the proof of Theorem 7 (a).

□

9.4 Proof of Theorem 7 (b)

Proof. Let \vec{G} be an orientation of $G = (V, E)$ of minimal Rényi entropy. Denote by $OPT = (q_i)_i$ the indegree distribution $q_i = \frac{v(i)}{W}$, where $v(i)$ is the sum of weights of all edges oriented in \vec{G} towards vertex $i \in V$, and W is the sum of all edge weights.

The Rényi entropy of OPT expands as follows:

$$\begin{aligned} H_\lambda(OPT) &= \frac{1}{1-\lambda} \log \sum_{i \in V} q_i^\lambda = \frac{1}{1-\lambda} \log \sum_{i \in V} \frac{v(i)}{W} \left[\frac{v(i)}{W} \right]^{\lambda-1} \\ &= \frac{1}{1-\lambda} \log \sum_{(i,j) \in \vec{G}} \frac{w_{i,j}}{W} \left[\frac{v(i)}{W} \right]^{\lambda-1} \end{aligned}$$

Since $x^{\lambda-1}$ is decreasing for $0 \leq \lambda < 1$ we infer

$$H_\lambda(OPT) \geq \frac{1}{1-\lambda} \log \sum_{(i,j) \in \vec{G}} \frac{w_{i,j}}{W} \left[\frac{\max\{v(i), v(j)\}}{W} \right]^{\lambda-1}$$

The inequality is true for any $\lambda > 1$ as well, as $x^{\lambda-1}$ is now increasing but we multiply with negative constant $\frac{1}{1-\lambda}$.

Let G^b be a biased orientation. Thus

$$v_{G^b}(i) = \sum_{(i,j) \in E, v(i) > v(j)} w_{i,j}$$

Let $BI = (q_i^b)_i$ be its indegree distribution. By the definition of biasedness we have:

$$\begin{aligned} H_\lambda(OPT) &\geq \frac{1}{1-\lambda} \log \sum_{(i,j) \in E} \frac{w_{i,j}}{W} \left[\frac{\max\{v(i), v(j)\}}{W} \right]^{\lambda-1} = \\ &= \frac{1}{1-\lambda} \log \sum_{(i,j) \in G^b} \frac{w_{i,j}}{W} \left[\frac{v(i)}{W} \right]^{\lambda-1} = \\ &= \frac{1}{1-\lambda} \log \sum_{i \in V} \frac{v_{G^b}(i)}{W} \left[\frac{v(i)}{W} \right]^{\lambda-1} = \frac{1}{1-\lambda} \log \sum_{i \in V} q_i^b \left[\frac{v(i)}{W} \right]^{\lambda-1} \end{aligned}$$

The Shapley value of an IS-game is [21] $s(i) = \frac{1}{2} \sum_{i \neq j} w_{i,j}$. Thus, the Shapley distribution $Sh = (s_i)_i$ of such a game is $s_i = \frac{s(i)}{W}$. Hence $H_\lambda(OPT) \geq$

$$\frac{1}{1-\lambda} \log \sum_{i \in V} q_i^b \left[\frac{2W \cdot s_i}{W} \right]^{\lambda-1} = \frac{1}{1-\lambda} \log \sum_{i \in V} q_i^b (s_i)^{\lambda-1} - 1$$

The difference between entropies of the optimal and Shapley distribution can be written as follows:

$$\begin{aligned} H_\lambda(OPT) - H_\lambda(Sh) &\geq \frac{1}{1-\lambda} \log \sum_{i \in V} q_i^b s_i^{\lambda-1} - 1 - \\ &- \frac{1}{1-\lambda} \log \sum_{i \in V} s_i^\lambda = \left[\frac{1}{1-\lambda} \log \sum_{i \in V} q_i^b (s_i)^{\lambda-1} + \right. \\ &+ \frac{1}{\lambda} \log \sum_{i \in V} s_i^\lambda - \frac{1}{\lambda(1-\lambda)} \log \sum_{i \in V} (q_i^b)^\lambda \left. \right] \\ &+ \frac{1}{\lambda(1-\lambda)} \log \sum_{i \in V} (q_i^b)^\lambda - \frac{1}{\lambda(1-\lambda)} \log \sum_{i \in V} s_i^\lambda - 1 \\ &= h_\lambda[BI, Sh] + \frac{1}{\lambda} H_\lambda(BI) - \frac{1}{\lambda} H_\lambda(Sh) - 1 \end{aligned}$$

Applying the discrete Gibbs Lemma we infer

$$H_\lambda(Sh) - H_\lambda(OPT) \leq \frac{1}{\lambda} (H_\lambda(Sh) - H_\lambda(BI)) + 1.$$

completing the proof. \square

Conclusions

The main contribution of this paper was to propose a parametric family of measures of worst-case fairness in cooperative settings. It raises many open questions, e.g. (a). obtain tight upper bounds for approximating strictly egalitarian, marginalistic and other WCF measures in IS games (b). obtain approximation guarantees for (general) TCSG games. (c). study WCF in other settings, e.g. NTU games such as *coalitional resource games* [22], or in the context of multiple coalitions. (d). study tradeoffs between fairness and other features of cooperative games.

References

- [1] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *Proceedings of the 18th annual ACM-SIAM SODA*, pages 189–198. Society for Industrial and Applied Mathematics, 2007.
- [2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [3] J. Arin, J. Kuipers, and D. Vermeulen. An axiomatic approach to egalitarianism in TU-games. *International Journal of Game Theory*, 37(4):565–580, 2008.
- [4] S. Arora and B. Barak. *Computational Complexity: A modern Approach*. Cambridge University Press, 2009.
- [5] Y. Bachrach and E. Elkind. Divide and conquer: False-name manipulations in weighted voting games. In *Proceedings of AAMAS’08*, pages 975–982, 2008.
- [6] Y. Bachrach, E. Elkind, R. Meir, D. Pasechnik, M. Zuckerman, J. Rothe, and J. S. Rosenschein. The cost of stability in coalitional games. In *Proceedings of SAGT’09*, pages 122–134. Springer, 2009.
- [7] Y. Bachrach, I. Kash, and N. Shah. Agent failures in totally balanced games and convex games. In *Proceedings of WINE’12*, pages 15–29. Springer, 2012.
- [8] Y. Bachrach, R. Meir, M. Feldman, and M. Tennenholtz. Solving cooperative reliability games. *Proceedings of UAI’11*, 2011.
- [9] Y. Bachrach, D. C. Parkes, and J. S. Rosenschein. Computing cooperative solution concepts in coalitional skill games. *Artificial Intelligence*, 204:1–21, 2013.
- [10] B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof Nash equilibria: I. concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.
- [11] B. D. Bernheim and M. D. Whinston. Coalition-proof Nash equilibria II. applications. *Journal of Economic Theory*, 42(1):13–29, 1987.

- [12] J. Cardinal and C. Dumeunier. Set covering problems with general objective functions. Technical Report 0802.2184, arXiv.org, 2008.
- [13] J. Cardinal, S. Fiorini, and G. Joret. Minimum entropy orientations. *Operations Research Letters*, 36(6):680–683, 2008.
- [14] J. Cardinal, S. Fiorini, and G. Joret. Minimum entropy combinatorial optimization problems. *Theory of Computing Systems*, 51(1):4–21, 2012.
- [15] G. Chalkiadakis, E. Elkind, M. Polukarov, and N. R. Jennings. The price of democracy in coalition formation. In *Proceedings of AAMAS’09*, pages 401–408, 2009.
- [16] G. Chalkiadakis, E. Elkind, and M. Wooldridge. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 2011.
- [17] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th ACM STOC*, pages 67–73, 2005.
- [18] T. Cover and J. Thomas. *Elements of information theory*. Wiley, 1991.
- [19] F. Cowell and K. Kuga. Inequality measurement: An axiomatic approach. *European Economic Review*, 15(3):287 – 305, 1981.
- [20] X. Deng, T. Ibaraki, and H. Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research*, pages 751–766, 1999.
- [21] X. Deng and C. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, pages 257–266, 1994.
- [22] P. E. Dunne, S. Kraus, E. Manisterski, and M. Wooldridge. Solving coalitional resource games. *Artificial Intelligence*, 174(1):20–50, 2010.
- [23] B. Dutta and D. Ray. A concept of egalitarianism under participation constraints. *Econometrica: Journal of the Econometric Society*, pages 615–635, 1989.

- [24] E. Elkind, L. A. Goldberg, P. Goldberg, and M. Wooldridge. Computational complexity of weighted threshold games. In *Proc. AAAI'07*, volume 22, page 718, 2007.
- [25] E. Halperin and R. Karp. The minimum entropy set cover problem. *Theoretical Computer Science*, 348(2–3):340–350, 2005.
- [26] A. Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proceedings of the 38th STOC*, pages 89–98. ACM, 2006.
- [27] J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr, and H. Gintis. *Foundations of human sociality: ethnography and experiments in fifteen small-scale societies*, Oxford University Press, 2004.
- [28] N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms*, 4(2):24, 2008.
- [29] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of STACS'99*, pages 404–413. Springer, 1999.
- [30] E. Lutwak, D. Yang, and G. Zhang. Cramér-Rao and moment-entropy inequalities for Rényi entropy and generalized Fisher information. *IEEE Transactions on Information Theory*, 51(2):473–478, 2005.
- [31] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
- [32] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [33] A. Roth, editor. *The Shapley Value: essays in honor of Lloyd S. Shapley*. Cambridge University Press, 1988.
- [34] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. M.I.T. Press, 2005.
- [35] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

- [36] T. Roughgarden and E. Tardos. Introduction to the inefficiency of equilibria. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [37] H. Theil. *Economics and information theory*. North-Holland, 1967.
- [38] D. Topkis. *Supermodularity and complementarity*. Princeton University Press, 1998.
- [39] A. Ullah. Entropy, divergence and distance measures with econometric applications. *Journal of Statistical Planning and Inference*, 49(1):137–162, 1996.
- [40] D. Williamson and D. Shmoys. *The design of approximation algorithms*. Cambridge University Press, 2011.
- [41] Y. Zick, M. Polukarov, and N. R. Jennings. Taxation and stability in cooperative games. In *Proceedings of the AAMAS'13 i*, pages 523–530, 2013.
- [42] M. Zuckerman, P. Faliszewski, Y. Bachrach, and E. Elkind. Manipulating the quota in weighted voting games. In *Proceedings of AAAI'08*, pages 215–220, 2008.