Hedonic Utility Games

Athina Georgara*
ageorg@iiia.csic.es
Artificial Intelligence Research Institute,
Spanish National Research Council
Barcelona, Spain

Georgios Chalkiadakis
gehalk@intelligence.tuc.gr
School of Electrical and Computer Engineering,
Technical University of Crete
Chania, Greece

ABSTRACT

We initiate the study of a novel class of cooperative games, the Hedonic Utility Games (HUGs), that takes into consideration both hedonic and utility-related preferences. We first formally define HUGs, and show how to extend and apply existing stability solution concepts to them. Then, we put forward the novel Individually Rational - Individually Stable (IRIS) solution concept, developed specifically for HUGs, that characterizes the stability of coalition structures in such settings. In addition, we propose a natural, "trichotomous" hedonic preferences model; study certain HUGs' properties in that model; and exploit it to characterize the feasibility of HUGs coalitions, and to obtain a probability bound for pruning the coalitional space, thus reducing the computational load of computing kernel-stable payoff configurations for IRIS partitions.

KEYWORDS

Game Theory, Cooperative Games, Hedonic Games, Multi-agent Systems

ACM Reference Format:
Athina Georgara and Georgios Chalkiadakis. 2020. Hedonic Utility Games. In *11th Hellenic Conference on Artificial Intelligence (SETN 2020), September 2–4, 2020, Athens, Greece.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3411408.3411412

1 INTRODUCTION

Cooperative games [9] can be naturally distinguished into utility-driven games and hedonic games. This reflects agents' motivation during coalition (group) formation. In utility-driven games[1], an agent seeks to acquire the best possible payoff, and therefore joins the coalition that offers her the highest reward. By contrast, in hedonic games [5], each agent is interested to participate in her most preferable coalition. One could say that in the former case the preference relation over coalitions is based on payoffs, while in the latter it is based on coalitional composition. However, in many real life settings, such an absolute demarcation among motives does not exist. On the contrary, people value (maybe in different proportions each) *both* hedonic preferences and payoff shares, when they are to collaborate with others in order to carry out a task. In the general case, when people are to form coalitions, they take all such motivating aspects into account.

For example, consider the way a startup company is created. In the early stages, there is a group of people sharing common ideas and perspectives, most likely a "core" initial group of friends. At the same time, these people value the potential economic profit the company's projects will yield. This is exactly the kind of settings we intend to formally describe and study in this work.

As such, in Section 2 we discuss a generic model that combines hedonic and utility aspects, and we introduce *hedonic utility games* (HUGs). In Section 3 we discuss the application of existing stability concepts into the HUGs setting, while in Section 4 we put forward our novel theoretical IRIS solution concept, and study its complexity. In Section 5 we first extend the well-known dichotomous hedonic preferences model to a natural trichotomous preferences one, and study HUGs and IRIS in that setting. As part of our work there, we characterise feasible coalitions for HUGs, and exploit this feasibility concept to obtain a probability bound for pruning the coalitional space that ultimately reduce the computational load for obtaining kernel-stable payoff configurations in IRIS partitions.

2 A GENERIC MODEL FOR HUGS

As mentioned, in this work we combine hedonic with utility games. That is, players have hedonic preferences over coalitions, but also form utility-based preferences over coalitions. The *hedonic preferences* take into consideration the identity of coalition members, i.e. each player only cares about which players are in her own coalition. The *utility-based preferences* are derived from a utility (characteristic) function and/or the payoff share each player receives. As such, a *hedonic utility game* in its generality is driven by two main components: the *hedonic aspect*, i.e. hedonic preferences over coalitions based solely on each coalition's members composition; and the *utility aspect*, i.e. the utility obtained by a coalition, that eventually leads to a payoff rewarded to each player. In other words, the hedonic preferences are the component that attributes a personalized opinion on a given coalition, while the utility function along with the utility-based preferences are the component that attributes a "generally accepted" quantified opinion on the same coalition.

Definition 2.1. (Hedonic Preferences) Let $N = \{1, \cdots, n\}$ be a finite set of players, a hedonic preference relation is denoted by $\succsim^{\text{hed}} = (\succsim_1^{\text{hed}}, \cdots, \succsim_n^{\text{hed}})$, where $\succsim_i^{\text{hed}} \subseteq N_i \times N_i$ is complete, reflexive, and transitive relation, with $N_i \subseteq 2^N$ and specifically $N_i = \{S \subseteq N \text{ such that } i \in S\}$.

Definition 2.2. (Utility-based Preferences) Let $N = \{1, \cdots, n\}$ be a finite set of players and $v : 2^N \to \mathbb{R}$ be a utility function that

*Supported by projects AI4EU (H2020-825619), 2019DI17, and *Enzyme Advising Group.*
[1]With this we refer to games such as TU/NTU Games, CFGs, etc. [9].

SETN 2020, September 2–4, 2020, Athens, Greece

ACM ISBN 978-1-4503-8878-8/20/09…$15.00
https://doi.org/10.1145/3411408.3411412

give rise to utility-driven preference relation $\succsim^{\text{ut}} = (\succsim_1^{\text{ut}}, \cdots, \succsim_n^{\text{ut}})$, where $\succsim_i^{\text{ut}} \subseteq N_i \times N_i$ is a complete, reflexive and transitive relation, with $N_i = \{S \subseteq N \text{ such that } i \in S\}$.

Note that in utility-driven games, the utility function denotes the estimated 'worth' of the coalition as a whole, which however is about to be distributed to the coalition members according to some payoff vector. So, even though in the above definition we do not explicitly refer to payoff vectors, clearly they can affect the utility-driven preferences as well; indeed, payoff allocation vectors are key in determining the outcomes of cooperative games, and are explicitely take into account in stability concepts such as the core [9]. That is, since a payoff vector is a fragmented version of the formed coalitions' utilities, the two notions are interrelated.

Thus, the player's utility-based preferences $\succsim_i^{\text{ut}}$ can be defined in various ways given the problem at hand. For instance, let v be a utility function, and $\mathbf{x}(C)$ a payoff vector for some coalition $C \subseteq N$, then, $\succsim_i^{\text{ut}}$ could be of the form $C_1 \succsim_i^{\text{ut}} C_2$ if $x_i(C_1) \geq x_i(C_2)$; or, $\succsim_i^{\text{ut}}$ could solely depend on the utility function: $C_1 \succsim_i^{\text{ut}} C_2$ if $v(C_1) \geq v(C_2)$. One step further, we can say that each player forms an overall ordering over coalitions that takes into account both hedonic and utility-based preferences.

Definition 2.3. (Overall Preferences) Let $\succsim_i^{\text{hed}}$ be a hedonic preference relation, and $\succsim_i^{\text{ut}}$ be a utility-based preferences relation, for some player i. Then, a function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}}) = \succsim_i^{\text{overall}}$, blends the hedonic and utility-based preferences in order to produce a single overall ordinal preference relation over coalitions.

Thus, a hedonic utility game in its generality is defined as follows:

Definition 2.4. (HUGs) A Hedonic Utility Game (HUG) $\mathcal{G}$ is given by a tuple $\langle N; h_1, \cdots, h_n; \succsim_1^{\text{hed}}, \cdots, \succsim_n^{\text{hed}}; v\rangle$, where for each $i \in N$ $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ is a blending function that produces an overall preference relation $\succsim_i^{\text{overall}}$ combined by the hedonic preference $\succsim_i^{\text{hed}}$ and a utility-based preference $\succsim_i^{\text{ut}}$ derived from the utility function $v : 2^N \to \mathbb{R}$. The outcome of a HUG is a pair $\langle \mathcal{CS}, \mathbf{x}\rangle$, where $\mathcal{CS}$ is a coalition structure; and $\mathbf{x} \in \mathbb{R}^n$ is a payoff vector related to $\mathcal{CS}$.

For each player i, function h_i may differ; so, let us examine how h can be formed in some base-line scenarios:

(i) A player i's overall preferences that depend *only* on the *hedonic* preferences, $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}}) = \succsim_i^{\text{hed}}$; i.e., even if $v(S) < v(T)$ or $x_i(S) < x_i(T)$, player i still prefers coalition S over T.

(ii) For some player i, function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}}) = \succsim_i^{\text{ut}}$ depends *exclusively* on the *utility-based* preferences, that is player i's overall preferences is $S \succsim_i^{\text{overall}} T$ if and only if $v(S) \geq v(T)$ (or if and only if $x_i(S) \geq x_i(T)$). Or the i's overall preferences can be based on what we later call 'potentially individual rationality', i.e., $S \succ_i^{\text{overall}} T$ if and only if $v(S) \geq \sum_{j \in S} v(\{j\})$ and $v(T) < \sum_{k \in T} v(\{k\})$.

(iii) Another case is that of depending on both *hedonic* and *utility-based* preferences, i.e. let for some player i function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ be defined as "i prefers coalition S over coalition T if and only if $S \succsim_i^{\text{hed}} T$ and $v(S) \geq v(T) - \varepsilon_i$"; where ε_i is a threshold corresponding to an acceptable utility-loss, determined by i. Similarly could be the case of "i prefers coalition S over coalition T if and only if $S \succsim_i^{\text{hed}} T$ and $x_i(S) \geq x_i(T) - \varepsilon_i$", where ε_i now determines an acceptable payoff-loss.

(iv) A quite more complex scenario is that player i's function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ be defined as "i prefers coalition S over coalition T if and only if $(S \succsim_i^{\text{hed}} T$ and $v(S) \geq v(T) - \varepsilon_i^{max})$ *or* $(v(S) \geq v(T) + \varepsilon_i^{min}$ regardless of the hedonic relation of i on S and T)"; where ε_i^{max} defines a maximum acceptable utility-loss for player i in order to satisfy her hedonic preferences, and ε_i^{min} defines a minimum desirable utility-gain for player i in order to ignore her hedonic preferences. Respectively, function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ could be also defined as "i prefers coalition S over coalition T if and only if $\left(S \succsim_i^{\text{hed}} T\right.$ and $\left.x_i(S) \geq x_i(T) - \varepsilon_i^{max}\right)$ *or* $\left(x_i(S) \geq x_i(T) + \varepsilon_i^{min}\right.$ regardless of the hedonic relation of i on S and $\left.T\right)$"; where now ε_i^{max} and ε_i^{min} correspond to a maximum acceptable loss and a minimum desirable gain, respectively.

These are just some scenarios: in general $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ can take any form. Thus, each player can value her hedonic and utility-based preferences differently. Note that this blending function h_i may result in a preference relation that is *non-transitive*; in our view, this is a very interesting property that alters the way the concept of 'rationality' is perceived in settings where both the hedonic and the utility aspect affect the outcome. In other words, in real-life settings where we deal with interpersonal relations *and* imminent payoff, rationality stops being straightforward and adapts in such a complex environment.[2]

To help our discussion in the rest of the paper, we now provide an alternative definition equivalent to Definition 2.4, but which explicitly refers to the $\succsim_i^{\text{hed}}$ and $\succsim_i^{\text{ut}}$ aspects of the problem. As such, HUGs' definition now becomes:

Definition 2.5. (HUGs–alternative) A Hedonic Utility Game $\mathcal{G}$ is given by a tuple $\langle N; \succsim^{\text{hed}}; v\rangle$, where $\succsim^{\text{hed}} = (\succsim_1^{\text{hed}}, \cdots, \succsim_n^{\text{hed}})$ is a vector of hedonic preference relations, one for each i; and $v : 2^n \to \mathbb{R}$ is a utility function. The outcome of a HUG is a pair $\langle \mathcal{CS}, \mathbf{x}\rangle$, where $\mathcal{CS}$ is a coalition structure; and $\mathbf{x} \in \mathbb{R}^n$ is a payoff vector related to $\mathcal{CS}$.

To ease notation, from now on we use $\succsim_i$ to refer to $\succsim_i^{\text{hed}}$, unless explicitly stated otherwise. Also, henceforth, we use $S \succ_i T$ to denote that i strictly prefers coalition S to T; and $S \sim_i T$ to denote that i is indifferent between coalitions S, T. Moreover, in the rest of the paper we use the following notation: N is a set of players and $N_i = \{S \subseteq N : i \in S\}$ is the set of all coalitions that contains agent i; π is a partition (also referred to as coalition structure $\mathcal{CS}$) of N, while $\pi(i) \equiv S \in \pi : i \in S$, is the single coalition within π that contains agent i.

[2] A HUG could be reminiscent of Multiple Objective Games (MOGs) [22] where each player has a multi-dimensional utility vector, and needs to optimise every dimension. However, in HUGs the utility function is common to all agents, while there is also a distinct hedonic dimension in the preferences of each player. Moreover, as we explained above, the HUG blending function can be non-transitive, while in MOGs the focus is on properly efficient optimisation solutions with explicit and implicit convexity and transitivity assumptions.

2.1 Real-life examples and applications

Though there are classes of games that can sufficiently model real cooperative problems, these models so far ignore either their hedonic or their utility aspect. Here, we give a few examples of settings where both aspects co-exist.

For instance, consider an online platform that promotes startup companies' formation. Such companies have a core (a group of friends) that shares ideas, passion, way of thinking etc.; a group of friends that most likely will be sceptical about cooperating with individuals that are in rivalry. Despite that, since they constitute a working firm, this core of people also care about the prosperity of the company, i.e., they consider the company's profits. In this scenario, we can quite clearly see the potential of the HUG model to capture both the hedonic and the utility aspect. That is, let each startup company be viewed as an initial coalition, while the individual candidates to join the startups as singletons. The individuals exhibit their personal preferences over *(a)* the existing coalitions, based on their interest regarding the companies' line of work (in order to join some existing startup), and *(b)* other singletons based on common interests and vision (in order to form a new startup). The members of the companies on the other hand, exhibit their personal preferences over *(c)* the individuals based on personality and other characteristics (in order to hire), and *(d)* other companies based on overlapping areas of work (in order to begin a bilateral corporation). At the same time, each collaboration is characterised by the commonly accepted 'worth' of the upcoming profits.

We may also think of the *social ridesharing* problem. In ridesharing, a set of commuters forms coalitions and arrange one-time rides at short notice. The goal is to *(a)* transport all commuters to their destination; and at the same time *(b)* minimize their expenses. In [7], the authors adopt a cooperative game theoretic approach in order to tackle the problem considering only the utility aspect of the game, i.e. they satisfy the two objectives *(a)* and *(b)*. Nevertheless, in many cases a commuter i may prefer to *rideshare with her friends* even if this ride costs more than others. In other words, a commuter may be willing to pay more if she is to spend time with people she is having fun with. Yet, if for a commuter the total cost of a ride with friends is way larger than a ride with completely strangers or a ride with people that the commuter dislikes, the commuter will need to weigh her hedonic preferences against the cost of the trip. Here, the pesonalized opinion (expressed through preferences) and the commonly accepted opinion (expressed through cost) clearly influence the commuter's decision to join or not to a specific ride, and thus a model such HUG that capture both can be used.

We could also think of a recommender system that is used by a travel agency. The travel agency is interested in creating groups of travellers that will *(a)* have a good time during their holidays, *(b)* meet their constraints/desires, for example a maximum total cost. Indirectly, a travel agent will consider these two objectives when planning the holiday packages of her clients. For instance, consider a travel agency that offers discount coupons to the fans of two rival local football clubs; the travel agent will avoid placing in the same holiday package fans from rival clubs, while they willprefering to place together fans of the same club. At the same time, each holiday package is characterised by some commonly accepted utility expressed via the total expenses, the dates, the destination, the number of different activities, etc. Thus, the recommender system used by the agency should form groups of travellers that appear to prefer each other's company, and at the same time plan holiday packages which meet the travellers' budget and expectations.

The above are only a few examples, where we can clearly observe the co-existence of the hedonic and the utility aspect that affect individual's decisions. In these settings the use of the HUG model would allow us to have a "holistic" perspective of the problem at hand, by explicitly taking into account both hedonic preferences and utilities, and apply solution concepts as the ones described in the following sections.

3 APPLYING EXISTING SOLUTION CONCEPTS TO HUGS

In this section we discuss how several existing stability solution concepts can be applied on the HUG model. We focus on concepts from Hedonic Games and the TU Games literature, to approach HUGs from both aspects.

3.1 Individual Rationality

First we discuss the concept of *individual rationality (IR)*. Individual rationality is a notion one finds on both hedonic and TU games [9]. In hedonic games, a partition π is individually rational if for every agent $i \in N$ it holds that $\pi(i) \succsim_i \{i\}$. In words, π is IR if each i prefers its current coalition, $\pi(i)$, at least as much as the singleton $\{i\}$ [5, 8]. In TU games, a partition π with respect to a payoff vector $\mathbf{x}$, is individually rational if for every agent i it holds that $x_i \geq v(\{i\})$. In words, π is IR if each agent i receives a payoff that is at least as good as what she can earn on her own [9].

Now, given a hedonic utility game $\mathcal{G} = \langle N; \succsim; v \rangle$, a partition π of N is individually rational if π is IR in both hedonic and TU terms. That is, *wrt* a payoff vector $\mathbf{x}$ related to π, for every $i \in N$ it holds that $\pi(i) \succsim_i \{i\}$, and at the same time, $x_i \geq v\big(\{i\}\big)$.

3.2 Individual Stability

The concept of individual stability of coalition structures is a key notion in the hedonic games literature [5, 8, 11]. In an individually stable (IS) partition, no agent prefers to unilaterally deviate into a new coalition and, at the same time, is welcomed by this new coalition. IS-deviation in HUGs is defined exactly as the usual IS-deviation.

Definition 3.1. (IS-deviation in HUGs) In a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, given a partition π, an agent i can IS-deviate into a coalition $S \in \pi \cup \{\emptyset\}$ if it holds that $S \cup \{i\} \succ_i \pi(i)$ and for each $j \in S$ it holds that $S \cup \{i\} \succsim_j^S$.

A partition π is individually stable if no agent can IS-deviate. By its definition, individual stability includes individual rationality, i.e. if a partition satisfies individual stability then it satisfies individual rationality (in the hedonic sense) as well. For pure hedonic games, computing or even deciding the existence of IS partitions is NP-complete [6]; thus, it is NP-complete to find an individually stable partition in a HUG.

For TU games there is no 'individual stability' solution concept;[3] however, in a utility game an IS-deviation can be thought of as follows: given a partition π, an agent i can IS-deviate into $S \in \pi \cup \{\emptyset\}$ iff $x_i(S \cup \{i\}) > x_i(\pi(i))$ and for each $j \in S$ we have that $x_j(S \cup \{i\}) \geq x_j(S)$. Thus, we can provide an *enhanced* definition for IS-deviation for HUGs as:

Definition 3.2. (Enhanced IS-deviation in HUGs) In a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, given a partition π, an agent i can IS-deviate into a coalition $S \in \pi \cup \{\emptyset\}$ if it holds that $S \cup \{i\} \succ_i \{i\}$ and $x_i(S \cup \{i\}) > x_i(\pi(i))$; and also, for each $j \in S$ it holds that $S \cup \{i\} \succsim_j S$ and $x_j(S \cup \{i\}) \geq x_j(S)$.

Note that we said nothing on how to compute the payoff $x_i(S)$; one can arbitrarily set a complete representation $\mathbf{X} : 2^n \to \mathbb{R}^n$, such that for every coalition $S \subseteq N$ there is a payoff $x_i(S) \in \mathbf{X}$ for each $i \in N$.[4]

3.3 Core Stability

The strongest cooperative solution concept regarding stability is the *core*, the set of outcomes where no subset of players has an incentive to deviate. The core solution concept is well-defined and well-studied in both hedonic games [5] and utility-driven games [9]. For interest and completeness, we define the core of HUGs in a straightforward manner.

Definition 3.3. (Core of HUGs) Given a hedonic utility game $\mathcal{G} = \langle N; \succsim; v \rangle$, a pair $\langle S, \mathbf{y} \rangle$ blocks the outcome $\langle \pi, \mathbf{x} \rangle$ if for every $i \in S$ holds $S \succ_i \pi(i)$ and $y_i > x_i$. The core $C(\mathcal{G})$ of a HUG is the set of all partitions π that admits no blocking pairs $\langle S, \mathbf{y} \rangle$.

3.4 Kernel Stability

The kernel consists of all outcomes where any pair of agents are in bilateral *equilibrium*; that is, no player can claim a share of another player's payoff [9, 10]. The kernel is always non-empty [18]. Nonetheless, computing the kernel is itself hard; Aumann, Peleg, and Robinowitz in [2] proposed a set of rules for determining the kernel, but this is impractical for large settings.

The Stearns transfer scheme [21] described a payoff transfer scheme that converges to the game's kernel. This transfer scheme performs a series of k-transfers that rearrange the payoff configuration such that each pair of agents in a coalition is in bilateral equilibrium. Regardless, this may require an infinite number of steps. Under that realisation, Shehory and Krauss in [19] provided a modification that allows fast convergence, given a specified error ϵ. Both schemes transform a payoff vector to a kernel-stable one with respect to some partition. Thus, in principle one could have a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, provided along with an IS coalition structure[5], and an arbitrary initial payoff vector $\mathbf{x}$, which can then be transformed into a kernel-stable one. However, this transformation procedure may result in a payoff configuration that is *neither individually stable* (under Enhanced IS-deviation), nor even *individually rational*; subsequently, that would allow incentives for deviations. We now introduce a new solution concept that helps us overcome this problem.

[3]One could, of course, consider the special case of the *core* [9] where an agent forms a profitable deviating coalition by joining an existing one or by staying alone.
[4]We make the intuitive assumption that for every $j \notin S$ we have $x_j(S) = 0$, and $x_i(\emptyset) = 0$ for each $i \in N$.
[5]Assuming that IS is non-empty.

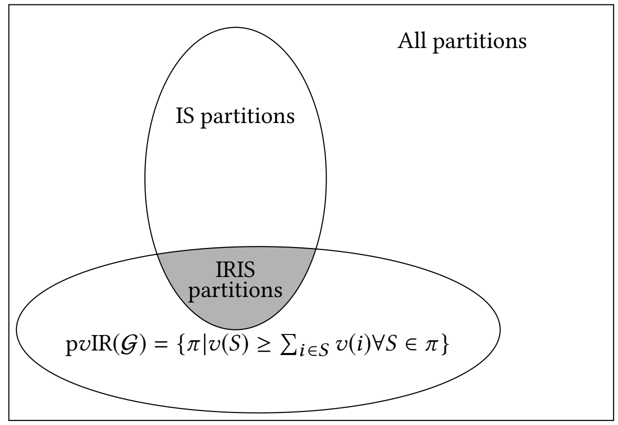

Figure 1: IRIS partition space

4 THE IRIS SOLUTION CONCEPT

In this section, we propose a novel solution concept designed specifically for the hybrid model of HUGs. In Subsection 3.2, we presented an extension of the notion of individual stability that considers utility as well. However, this enhanced version of individual stability for HUGs, requires an explicitly 'predefined' payoff configuration space, a payoff configuration per each of all possible partitions. That would be prohibitive, in terms of space, even for a small, let alone for a large number of players. Here, we propose the novel *individually rational - individually stable (IRIS)* solution concept that does not suffer from this problem. In IRIS, we have two requirements:

- the partition is individually stable as far as the hedonic aspect of the game is concerned, and
- the partition is such that an individually rational payoff *can potentially* be provided.

In other words, we seek partitions that satisfy the hedonic individual stability concept, and at the same time the coalitional values are such that all players can claim a payoff that is at least as good as what they can earn on their own; i.e., partitions that are hedonically individually stable, and partition that provide *imputations* (for instance, kernel-stable partitions). Such (IRIS) partitions (may) exist in the intersection of the set of individually stable partitions and the set of partitions where each agent can at least claim what she can earn on her own $\{\pi : v(S) \geq \sum_{i \in S} v(\{i\}), \forall S \in \pi\}$ as illustrated in Figure 1. Let us denote the set of individual stable partitions as IS($\mathcal{G}$), and the set $\{\pi : v(S) \geq \sum_{i \in S} v(\{i\}), \forall S \in \pi\}$ as (**p**otentially v-**I**ndividually **R**ational) pvIR($\mathcal{G}$). To formally define IRIS, let us first define the concept of *v-rationalizing deviation*:

Definition 4.1. (v-Rationalizing Deviation) Given a partition π of N, a deviation of agent i from $T = \pi(i)$ into $S \in \pi \cup \{\emptyset\}$ is called *v-rationalizing* if $v(T) < \sum_{j \in T} v(\{j\})$ and $v(S \cup \{i\}) \geq v(\{i\}) + \sum_{k \in S} v(\{k\})$.

An agent i that has an incentive for v-rationalising deviation from coalition T to coalition S, has in fact a blending function h_i that produces an overall preference relation depending only on the utility function v as follows: "agent i strictly prefers any coalition S such that $v(S \cup \{i\}) \geq v(\{i\}) + \sum_{k \in S} v(\{k\})$ to any coalition T such that $v(T) < \sum_{j \in T} v(\{j\})$, i.e. $S \succ_i^{\text{overall}} T$, and is indifferent in any other case".

We can now exploit Definition 3.1 and Definition 4.1 to construct the formal expression of *IRIS deviation*, and the characterisation of *IRIS partitions*:

Definition 4.2. (IRIS-deviation) Given a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ and a partition π, agent i can *IRIS-deviate* into $S \in \pi \cup \{\emptyset\}$if

- i can IS-deviate into S; **or**
- i can perform a *v-rationalizing* deviation into S.

Definition 4.3. (IRIS partition) Given a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, a partition π is *individually rational-individually stable* if no agent can IRIS-deviate. That is,

(1) $\nexists i, S \in \pi \cup \{\emptyset\}$ s.t. $S \cup \{i\} \succ_i \pi(i)$, and $S \cup \{i\} \succsim_j S, \forall j \in S$; **and**
(2) $\forall i \in N$ it holds that $v(\pi(i)) \geq \sum_{k \in \pi(i)} v(\{k\})$.

The second condition in Definition 4.3 ensures that in an IRIS partition, there *can* exist a payoff configuration that is individually rational. In utility-driven games the payoff vector results from the distribution of each coalition's utility to its members. In IRIS partitions, by allowing *only* the coalitions that can afford to reward their members with a payoff at least as good as their individual utility, we can guarantee that there exists at least one payoff configuration that is individually rational. We clarify that we slightly abuse the term 'individual rationality', since we do not consider any payoff vector in particular, but coalition structures that can *potentially lead* to an individually rational payoff configuration.

Under this concept the agents that have motive to IRIS-deviate follow the blending function $h_i(\succsim_i^{\text{hed}}, \succsim_i^{\text{ut}})$ defined as "agent i strictly prefers coalition S to T, $S \succ_i^{\text{overall}}$, if and only if $S \succ_i^{\text{hed}} T$ and $\forall\, j \in S\; S \succsim_j^{\text{hed}} S \setminus \{i\}$, **or** $v(S \cup \{i\}) \geq v(\{i\}) + \sum_{j \in S} v(\{j\})$ and $v(T) < \sum_{j \in T} v(\{j\})$".

Note that we are able to define and use the IRIS solution concept exactly due to the hybrid nature of HUG settings. That is, we could *not* have had the IRIS solution concept in a pure utility-driven game, since the notion of individual stability is not defined in such settings; nor could we have had it in pure hedonic games where the notion of utility is not defined. Notice, however, that by dropping the first condition of Definition 4.3 we would be able to consider a special case of individual rationality in pure utility-driven settings where condition (2) holds. At the same time, by dropping the second condition we would end up with the concept of individual stability in pure hedonic games; as such, Definition 4.3 generalizes the individual stability concept to HUGs. Notice also that IRIS is a strengthening of individual stability: every IRIS partition is always IS, but the opposite is not necessarily true.

Complexity of IRIS. In Algorithm 1 we provide an $O(n)$ algorithm that checks if an agent can IRIS-deviate into a coalition.

Proposition 4.4. *The problem* isInIRIS *(Fig. 2) is decidable in polynomial time.*

Proof. Go through Algorithm 2. The conjunction of the for-loops in lines 1 and 2 executes at most n^2 times, while for checking CanIRIS-Deviate(i,S,$\succsim$,v,π) we need at most $3 \cdot n$ computations ($O(n)$). Thus we can check in polynomial time, $O(n^3)$, whether a given partition of a HUG is IRIS or not. □

Name: isInIRIS
Instance: A HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, and a partition π
Question: Is π in IRIS($\mathcal{G}$)?

Name: ExistsIRISPartition
Instance: A hedonic utility game $\mathcal{G} = \langle N; \succsim; v \rangle$
Question: Is there a π that is an IRIS partition?

Figure 2: Two IRIS-related decision problems.

Algorithm 1: CanIRIS-Deviate(i,S,$\succsim$,v,π)

```
if( CanIS-Deviate(i,S,≿,π) ): return True;
current_value ← ∑_{k∈π(i)} v(k);
if( current_value ≥ v(π(i)) ): return False;
new_value ← ∑_{j∈S} v(j) + v(i);
if( new_value ≥ v(S ∪ {i}) ): return True;
return False;
```

Algorithm 2: checkIRISPartition($\mathcal{G}$,π)

```
1 for every existing coalition S ∈ π do
2     for every agent i ∈ N \ S do
3         if( CanIRIS-Deviate(i,S,≿,v,π) ): return False;
4 return True;
```

Even though it is easy to decide if a given partition is IRIS, it is also essential to solve the decision problem ExistsIRISPartition (see Figure 2). To answer this problem we need to either find a partition that satisfies conditions (1) and (2) of Definition 4.3, or decide that there is no such partition. However, in general, these two conditions are completely unrelated; that is, since they refer to different aspects of a HUG (the former to the hedonic aspect, while the latter to the utility aspect), having information that regards the first condition, provides us with no information regarding the second one, and vice versa. Therefore, a machine that decides the ExistsIRISPartition problem needs to solve *two unrelated, separate problems*, and come with a partition that is an answer to both problems or *halt* if there is none such partition. Since one of the problems is NP-complete [6], ExistsIRISPartition is NP-hard. We can also explicitly prove Proposition 4.5 below:

Proposition 4.5. *It is NP-hard to find an IRIS partition in a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ with arbitrary hedonic preferences.*

Proof. Suppose we have a hedonic game $\langle N, \succsim \rangle$ with arbitrary preference relations, exactly as the model considered in [6]. Add to this game a superadditive utility function v to get a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$. Due to superadditivity, condition 2 of Definition 4.3 always stands; thus a partition π is IRIS if and only if it satisfies condition 1 of Definition 4.3, i.e. if and only if it is IS. However, Ballester showed in [6] that it is NP-complete to find an IS partition in a game with arbitrary hedonic preferences. Therefore, it is NP-hard to find an IRIS partition with arbitrary preferences. □

4.1 A randomized transition scheme to reach an IRIS partition

In Proposition 4.5 we showed that it is NP-hard to build an IRIS partition. However, suppose we have an oracle which informs us that IRIS($\mathcal{G}$) is non-empty for some HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, i.e. that there is at least one partition π is in IRIS, and more specifically we are interested in a specific such π.[6] In this case, we can use a randomized transition scheme to reach a point in IRIS($\mathcal{G}$).

We begin with a random partition; if this partition is the IRIS partition of interest, we stop. Otherwise, we successively perform IRIS deviations. In case we get stuck in a loop, i.e., the partition transitions from point A to point B and vice versa (or if we reach an IRIS partition different to the one of interest), we perform a random, not necessarily IRIS, deviation. At some point, following this procedure we will reach the IRIS partition, however this may take an arbitrarily large number of steps. Nonetheless, given that we know an IRIS partition exists, the time needed for the transition scheme to converge into an IRIS partition is highly dependent on the sizes of the IS($\mathcal{G}$) and pvIR($\mathcal{G}$) sets, and the proportion of their intersection with respect to the total area of partitions (see Figure 1). In real-world settings with IRIS partitions, we anticipate that this time will not be prohibitive in practice (intuitively, the IRIS intersection area would be either large or very small with respect to the IS(G) and pvIR(G) ones). Of course, this has to be verified in specific environments.

5 INSTANTIATION OF HUGS TO TRICHOTOMOUS PREFERENCES

So far within the HUGs framework, we considered the hedonic preferences to be arbitrary. However, here we present a modification of the well-known *dichotomous hedonic preferences* that will allow us to obtain certain algorithmic results for solution concepts in HUGs. Dichotomous preferences were introduced in [4] to suggest that for each player the related coalition space can be partitioned into two disjoint subsets N_i^+ and N_i^-. In our trichotomous preferences modification we now explicitly require that each player i:

- strictly prefers all coalitions in N_i^+ to singleton,
- strictly prefers singleton to all coalitions in N_i^-, and
- is indifferent about coalitions in the same subset.

That is, for some $S, T \in N_i$ we have $S \succ_i \{i\} \succ_i T$ if and only if $S \in N_i^+$ and $T \in N_i^-$, and $S \sim_i T$ if and only if $S, T \in N_i^+$ or $S, T \in N_i^-$. Also, $\forall i \; N_i = N_i^+ \cup \{i\} \cup N_i^-$. The trichotomous preferences model is an intuitive paradigm that holds in many real life settings; for instance, consider work groups for a school project assignment, a student would prefer to be in a group with her friends than being alone, and in the same time would prefer being alone rather in group of people she dislikes. Therefore, if we adopt trichotomous preferences, we can build on [17] and obtain an IS partition in $O(n^3)$ using Algorithm 3.

[6]Assume that π is, e.g., one of interest to an "oracle" external entity which, however, cannot enforce its creation. In other words, this external entity is aware of a specific partition π that it is in IRIS($\mathcal{G}$), and can only reveal its existence and confirm whether a partition is the desired one.

Algorithm 3: IS Partition(N,$\succsim$,v)

$\pi \leftarrow \emptyset$;
$\forall i \in N$ assign i into $\{i\}$;
while *there are agents that can IS-deviate* **do**
  for *every agent i that is singleton* **do**
    for *every existing coalition $S \in \pi$* **do**
      if(*CanIS-Deviate(i,S,$\succsim$,π)* **):**
        assign agent i into coalition S;
return π;

Algorithm 4: CanIS-Deviate(i,S,$\succsim$,π)

if(*$S \cup \{i\}$ is strictly preferable to i than her current coalition $\pi(i)$* **):**
  for *every agent $j \in S$* **do**
    if(*S is strictly preferable to j than coalition $S \cup \{i\}$* **):**
      /* i cannot IS-deviate into S since j objects this deviation. */
      return *False*
  /* At this point there is no objection, thus i can IS-deviate into S. */
  return *True*
return *False*

Having in mind, that for HUGs with trichotomous preferences the IS($\mathcal{G}$) is non-empty, and that we can reach an IS partition in polynomial time, along with the fact that the pvIR($\mathcal{G}$) is always non-empty for any HUG, the immediate question arises: for a HUG with trichotomous hedonic preferences, is the intersection IRIS($\mathcal{G}$) = IS($\mathcal{G}$) ∩ pvIR($\mathcal{G}$) non-empty? The answer is we have no guarantees that even under these simplifying assumptions an IRIS partition exists. For instance consider example 5.1.

Example 5.1. Consider a 4-player HUG with trichotomous preferences such that: $N_1^+ = \big\{\{1,2\},\{1,3,4\},\{1,2,4\}\big\}$, $N_2^+ = \big\{\{2,4\},\{1,2,4\},\{1,2,3,4\}\big\}$, $N_3^+ = \big\{\{1,3,4\},\{1,3\}\big\}$, $N_4^+ = \big\{\{1,4\},\{2,4\},\{1,2,4\},\{1,3,4\}\big\}$, and $v(\{i\}) = 2 \; \forall \; i \in N, v(S) = |S| \; \forall \; S \subseteq N : \; |S| = 2, 3$, and $v(N) = 8$. According to these preferences, the individually stable partitions are the following:

$$\text{IS}(\mathcal{G}) = \Big\{\big\{\{1,2,4\},\{3\}\big\},\big\{\{1,3,4\},\{2\}\big\}\Big\}$$

while according to v the potentially individually rational partitions are:

$$\text{p}v\text{IR}(\mathcal{G}) = \Big\{\big\{\{1\},\{2\},\{4\},\{3\}\big\},\big\{\{1,2,3,4\}\big\}\Big\}.$$

Therefore, IRIS($\mathcal{G}$) ≡ IS($\mathcal{G}$) ∩ pvIR($\mathcal{G}$) = $\emptyset$, there is no IRIS partition even if the setting is a simplified HUG with trichotomous preferences.

Now, if the oracle verifies that for a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ with trichotomous preferences the IRIS($\mathcal{G}$) is non-empty, then we can use the aforementioned transition scheme (Subsection 4.1) and reach an IRIS partition where every coalition is individually stable in terms of hedonic preferences, and also has the 'capability' to provide an individually rational payoff vector. However, we said nothing so far about the agent's actual payoffs.

One could use the classic transfer scheme proposed by [19], and build a kernel-stable configuration for any IRIS coalition structure. Nonetheless, as pointed out by [7], the Shehory-Kraus transfer scheme will become eventually inefficient when the number of agents is increased. In [7] Bistaffa *et al.* overcome this problem by not considering infeasible coalitions. In [19] the authors also discuss about predefining an acceptable range of sizes for coalitions in order to achieve polynomial complexity. Similarly, in a HUG we can reduce the coalitional search space by disregarding infeasible coalitions. Although in general we may not be able to perform any pruning, under the assumption of trichotomous preferences we can discard infeasible coalitions.

Definition 5.2. (Feasibility via trichotomous properties) Given a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ with trichotmous preferences, a coalition $S \subseteq N$ is *infeasible* if and only if for at least one $i \in S$ it holds that $S \in N_i^-$, otherwise S is *feasible*:

$$S \text{ infeasible}^{\text{trich}} \Leftrightarrow \exists\, i \in S : S \in N_i^-,$$
$$S \text{ feasible}^{\text{trich}} \Leftrightarrow \forall i \in S : S \in N_i^+ \textbf{ or } |S| = 1.$$

According to Definition 5.2, we consider as feasible coalitions only those that are hedonically immune to deviations. That is, any coalition $S \in N_i^-$ for some $i \in S$, is unstable since agent i would prefer to deviate into a singleton. Thus, if a coalition is non-acceptable for at least one of its members, then this coalition is unstable, and therefore there is no benefit in being included in the computations of the kernel. Note that by disregarding infeasible coalitions we lose no individually stable partitions, i.e., we lose no coalitions that can be part of any individually stable partition:

Proposition 5.3. *(No loss of IS partitions by pruning) Given an IS partition π of a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ with trichotomous preferences, $\nexists\, S \in \pi$ s.t. $S \in N_i^-$ for some $i \in S$.*

Proof. The hedonic preferences in a HUG are defined according to trichotomous preferences model. That is for any agent $i \in N$ it holds that $S \succ_i \{i\} \succ_i T$ if and only if $S \in N_i^+$ and $T \in N_i^-$; in other words, each agent i strongly prefers to be member of any coalition in the set N_i^+ rather than being alone, but also she strongly prefers to be alone rather than being in any coalition in the set N_i^-. Now, assume there is a coalition $S \in \pi$ s.t. $S \in N_i^-$ for some $i \in S$. As indicated above, this particular agent i, strongly prefers to be on her own instead of being in S, motivating i to deviate into an empty coalition ($\{\}$). This means that π is not stable. *Therefore, in a individually stable HUG partition π, there is no coalition $S \in \pi$ s.t. $S \in N_i^-$ for some $i \in S$.* □

Proposition 5.4. *(No loss of IRIS partitions by pruning) Given an IRIS partition π of a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$ with trichotomous preferences, $\nexists\, S \in \pi$ s.t. $S \in N_i^-$ for some $i \in S$.*

Since IRIS partitions are IS, and pruning is not related to their pvIR($\mathcal{G}$) component, we *lose no IRIS partitions* either.

We may also consider the following simple setting: the trichotomous preferences over coalitions are actually lifted preferences over players. That is, let each $i \in N$ develops an "empathy" value e_i^j towards any other player $j \in N$, expressing i's perception as to how well it can collaborate with j; and a coalition $S \subseteq N \setminus \{i\}$ is placed in N_i^+ if a function $f(S, \mathbf{e}_i)$ meets a threshold t_i. This function f can be a summation of values e_i^j over the j in the coalition (as in Additively Separable Hedonic Games [5]), or average of these values (as in Fractional Hedonic Games [3]), or the pairwise average of these values: $f(S, \mathbf{e}) = \frac{1}{2} \cdot \sum_{i \in S} \sum_{j \in S} \frac{e_i^j + e_j^i}{2}$. Now, if we let each value $e_i^j \sim \mathcal{N}(\mu_i, \sigma_i^2)$ be Gaussian i.i.d. random variables, the pairwise average $f(S, \mathbf{e}) \sim \mathcal{N}(\frac{|S|}{2} \sum_{i \in S} \mu_i, \frac{|S|^2}{128} \sum_{i \in S} \sigma_i^2)$ is also a Gaussian random variable. Thus, using this universal (common to all agents) pairwise average function, we can obtain a probability bound on when a coalition is pruned according to trichotomous preferences.

Proposition 5.5. *(Prunning Probability Bound) Given a HUG $\mathcal{G} = \langle N; \succsim; v \rangle$, with trichotomous preferences following i.i.d. $e_i^j \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and the pairwise average function $f(S, \mathbf{e})$, a coalition $S \subseteq N$ is pruned with probability:*

$$P(S \text{ be trich-pruned}) \geq 1 - \frac{|S| \cdot \sum_{i \in S} \mu_i}{2 \cdot \max_{i \in S} t_i}.$$

Proof. Consider feasibility according to Def. 5.2; we could reform the condition as follows: S feasible$^{\text{trich}} \Leftrightarrow f(S, \mathbf{e}) \geq \max_{i \in S}\{t_i\}$. Thus, the probability of a coalition to be pruned exploiting the trichotomous preferences model, is:

$$P(S \text{ be trich-pruned }) = 1 - P(S \text{ feasible}^{\text{trich}}\,) = 1 - P\Big(f(S, \mathbf{e}) \geq \max_{i \in S}\{t_i\}\Big).$$

Now, exploiting Markov's Inequality [16] we have that:

$$\begin{aligned}
P\Big(f(S, \mathbf{e}) \geq \max_{i \in S}\{t_i\}\Big) &\leq \frac{E[f(S, \mathbf{e})]}{\max_{i \in S}\{t_i\}} \Leftrightarrow \\
P\Big(f(S, \mathbf{e}) \geq \max_{i \in S}\{t_i\}\Big) &\leq \frac{\frac{|S|}{2} \cdot \sum_{i \in S} \mu_i}{\max_{i \in S}\{t_i\}} \Leftrightarrow \\
1 - P\Big(f(S, \mathbf{e}) \geq \max_{i \in S}\{t_i\}\Big) &\geq 1 - \frac{|S| \cdot \sum_{i \in S} \mu_i}{2 \cdot \max_{i \in S}\{t_i\}} \Leftrightarrow \\
P(S \text{ be trich-pruned }) &\geq 1 - \frac{|S| \cdot \sum_{i \in S} \mu_i}{2 \cdot \max_{i \in S}\{t_i\}} \quad .
\end{aligned}$$

□

Kernel computation. As mentioned in Section 3.4, given a partition we can construct a payoff vector that lies in the kernel. Given an IRIS partition, we can in fact construct such a payoff vector in finite time, using an ϵ-kernel transfer scheme. Specifically, we can use the transfer scheme described in Algorithm *"Payments in the Kernel"* of [7] while considering *only* feasible coalitions characterised as in Definition 5.2, and thus compute a kernel stable payoff vector. (See Appendix A for the algorithm in question.)

6 CONCLUSIONS AND FUTURE WORK

In this paper we presented a novel hybrid class of cooperative games, HUGs, that couples hedonic preferences with utility ones. We extended several traditional stability concepts to the HUGs setting, via equipping them with the ability to cope with both utility and hedonic preferences; and proceeded to propose IRIS, a novel, HUGs-specific solution concept that combines the key notions of individual stability and individual rationality, and studied its computational properties. We then provided an instantiation of HUGs, along with a definition for characterising a coalition's feasibility, and used it to prune the coalitional space in order to compute kernel-stable payoffs for IRIS partitions. Last but not least, we provided a probability bound for pruning coalitions in HUGs.

We envisage this work to initiate the further study of this novel model. For instance, additional study of the HUGs' computational properties is required, along with a systematic evaluation of the time needed to converge to an existing point in IRIS($\mathcal{G}$). Due to the probable emptiness of IRIS, devising ways to identify *approximately* IRIS-stable partitions seems to be an interesting research direction. Another interesting perspective of future work is to investigate how a generic model such the one presented in [1] can be applied on overall preferences derived by blending functions (Def. 2.3) that may be non-transitive. Moreover, it would be interesting to examine how such a blending function can be extended in cooperative games in partition function form [14, 15]. Furthermore, we believe that characterising the optimality of and providing methods for the optimisation of non-transitive HUGs blending functions, perhaps in line with recent work on multiobjective optimisation with variable ordering structures [12, 13], is promising future work. Finally, it would be interesting to study HUGs under uncertainty, both from a practical machine learning viewpoint and from a theoretical one—by employing, e.g., PAC learning [20].

REFERENCES

[1] KRZYSZTOF R. APT and ANDREAS WITZEL. 2009. A GENERIC APPROACH TO COALITION FORMATION. *International Game Theory Review* 11, 03 (2009), 347–367. https://doi.org/10.1142/S0219198909002352 arXiv:https://doi.org/10.1142/S0219198909002352

[2] Robert J. Aumann, Bazalel Peleg, and Paul Robinowitz. 1965. A Method for Computing the Kernel of n-Person Games. *Mathematical Computation* 19 (1965), 531–551.

[3] Haris Aziz, Felix Brandt, and Paul Harrenstein. 2014. Fractional Hedonic Games. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems* (Paris, France) *(AAMAS '14)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 5–12.

[4] Haris Aziz, Paul Harrenstein, Jérôme Lang, and Michael Wooldridge. 2016. Boolean Hedonic Games. In *Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*. 166–175.

[5] Haris Aziz, Rahul Savani, and Hervé Moulin. 2016. Hedonic Games. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D.Editors Procaccia (Eds.). Cambridge University Press, 356–376.

[6] Coralio Ballester. 2004. NP-completeness in hedonic games. *Games and Economic Behavior* 49, 1 (2004), 1 – 30. https://doi.org/10.1016/j.geb.2003.10.003

[7] Filippo Bistaffa, Alessandro Farinelli, Georgios Chalkiadakis, and Sarvapali D. Ramchurn. 2017. A cooperative game-theoretic approach to the social ridesharing problem. *Artificial Intelligence* 246 (2017), 86 – 117.

[8] Anna Bogomolnaia and Matthew O. Jackson. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior* 38, 2 (2002), 201 – 230.

[9] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. *Computational Aspects of Cooperative Game Theory (Synthesis Lectures on Artificial Inetlligence and Machine Learning)* (1st ed.). Morgan & Claypool Publishers.

[10] Morton Davis and Michael Maschler. 1963. The kernel of a cooperative game. *Naval Research Logistics Quarterly* 12, 3 (1963), 223–259.

[11] Jacques H. Dréze and Jerald Greenberg. 1980. Hedonic Coalitions: Optimality and Stability. *Econometrica* 48, 4 (1980), 987–1003.

[12] Gabriele Eichfelder. 2014. Numerical Procedures in Multiobjective Optimization with Variable Ordering Structures. *Journal of Optimization Theory and Applications* 162, 2 (01 Aug 2014), 489–514. https://doi.org/10.1007/s10957-013-0267-y

[13] Gabriele Eichfelder and Tobias Gerlach. 2016. Characterization of properly optimal elements with variable ordering structures. *Optimization* 65, 3 (2016), 571–588. https://doi.org/10.1080/02331934.2015.1040793 arXiv:https://doi.org/10.1080/02331934.2015.1040793

[14] Athina Georgara, Dimitrios Troullinos, and Georgios Chalkiadakis. 2019. Extracting hidden preferences over partitions in hedonic cooperative games. In *Proc. of the 12th Knowledge Science Engineering Management (KSEM 19).*

[15] Tomasz P. Michalak, Andrew Dowell, Peter McBurney, and Michael Wooldridge. 2008. Optimal Coalition Structure Generation In Partition Function Games. 388–392. https://doi.org/10.3233/978-1-58603-891-5-388

[16] Michael Mitzenmacher and Eli Upfal. 2005. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis.* Cambridge University Press.

[17] Dominik Peters. 2016. Complexity of Hedonic Games with Dichotomous Preferences. In *AAAI.*

[18] David Schmeidler. 1969. The Nucleolus of a Characteristic Function Game. *SIAM J. Appl. Math.* 17, 6 (1969), 1163–1170.

[19] Onn Shehory and Sarit Kraus. 1999. Feasible Formation of Coalitions Among Autonomous Agents in Non-Super-Additive Environments.

[20] Jakub Sliwinski and Yair Zick. 2017. Learning Hedonic Games. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17.* 2730–2736.

[21] Richard E. Stearns. 1968. Convergent Transfer Schemes for N-Person Games. *Trans. Amer. Math. Soc.* 134, 3 (1968), 449–459.

[22] J. Zhao. 1991. The equilibria of a multiple objective game. *International Journal of Game Theory* 20, 2 (01 Jun 1991), 171–182.

A APPENDIX - KERNEL COMPUTATION IN IRIS PARTITION

With π^{IRIS} we denote an IRIS partition, v is a utility function, and ϵ is small positive number.

Algorithm 5: IRIS-ϵ-Kernel(π^{IRIS}, v, ϵ)

```
x = [ ];
for S ∈ π^IRIS do
    for i ∈ S do
        x_i = v(S)/|S|;
repeat
    /* Repeatedly compute surplus matrix (see Algorithm 6) and
       perform payoff transfers                                */
    s = ComputeSurplusMatrix(π, v, x);
    δ = max_{i,j∈S, S∈π^IRIS} {s_{i,j} − s_{j,i}};
    (i*, j*) = arg max_{i,j∈S, S∈π^IRIS} {s_{i,j} − s_{j,i}};
    if( x_{j*} − v({j*}) ≥ δ/2 ): d = x_{j*} − v({j*});
    else : d = δ/2;
    x_{j*} −= d; x_{i*} += d; // transfer from j* to i*
until ( δ / v(π^IRIS) ≤ ε );
return x;
```

Algorithm 6: ComputeSurplusMatrix($\pi, v, \mathbf{x}$)

```
s = −∞; // set all s_{i,j} = −∞
for every S ⊆ N : S feasible^trich do
    e_S = v(S) − Σ_{i∈S} x_i;
    for every i ∈ S do
        S_i = π(i); for j ∈ S_i \ S do
            s_{i,j} = max{s_{i,j}, e_S};
return s;
```