



HAL
open science

Nash Stability in Hedonic Skill Games

Laurent Gourvès, Gianpiero Monaco

► **To cite this version:**

Laurent Gourvès, Gianpiero Monaco. Nash Stability in Hedonic Skill Games. the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), May 2024, Auckland, New Zealand. hal-04626037

HAL Id: hal-04626037

<https://hal.science/hal-04626037v1>

Submitted on 26 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Nash Stability in Hedonic Skill Games

Laurent Gourvès¹ Gianpiero Monaco²

1. Université Paris Dauphine-PSL, CNRS, LAMSADE, France

2. University of Chieti-Pescara, Italy

laurent.gourves@dauphine.fr, gianpiero.monaco@unich.it

Abstract

This article deals with hedonic skill games, the strategic counterpart of coalitional skill games which model collaboration among entities through the abstract notions of tasks and the skills required to complete them. We show that deciding whether an instance of the game admits a Nash stable outcome is NP-complete in the weighted tasks setting. We then characterize the instances admitting a Nash stable outcome in the weighted tasks setting. This characterization relies on the fact that every agent holds (resp., every task requires) either a single skill or more than one skill. For these instances, the complexity of computing a Nash stable outcome is determined, together with the possibility that a natural dynamics converges to a Nash stable outcome from any initial configuration. Our study is completed with a thorough analysis of the price of anarchy of instances always admitting a Nash stable outcome.

Keywords

Hedonic Games, Nash Stable Outcomes, Price of Anarchy

1 Introduction

Coalitional skill games (CSGs, in short) model collaboration among entities and are based on the abstract notions of tasks and the skills required to complete them. Introduced in [BR08] (see also [BPR13]), CSGs are highly expressive cooperative games, but general CSGs are not concisely representable, meaning that their manipulation necessitates prohibitive calculations. Weighted Task Skill Games (WTSGs) are special CSGs which allow a succinct representation, yet preserving a good level of expressiveness [BR08, BPR13].

In WTSGs, we are given a set of agents, each one possessing a set of skills, and a set of weighted tasks, each one requiring a set of skills in order to be performed. An outcome for these games is a coalition structure, which is a partition of the agents into coalitions. A coalition can perform a task only if its members cover the set of required skills for the task. The worth of a coalition is defined as the sum of the weights of the tasks that the coalition can perform.

WTSGs have applications in many real life situations including oil extraction, voting, knowledge sharing, robots and human rescuers, and multi-sensor networks [BMJK10, BPR13].

So far, WTSGs are viewed from the angle of cooperative game theory. Typical fundamental questions in this field are the following: How to distribute the worth of collaboration that all participants find acceptable? Who are the powerful members of a coalition? Which coalition structure induces the maximum total value?

In this work, we consider WTSGs from the different perspective of hedonic games [DG80, AS16]. Our motivation is to capture and study situations where every agent can freely decide which coalition she wants to be part of, given that every agent's utility only depends on the members of her coalition. This different approach leads to new questions and challenges regarding WTSGs: Does the system admit a stable outcome? How difficult is the computation of such a state? Do the agents naturally converge to a stable state? How bad can the social welfare of stable states be?

The study of hedonic WTSGs (hedonic skill games, in short) is relevant in voting situations where the agents are candidates, the skill set consists of opinions on societal issues, and the tasks represent some segments of the electorate that are sensitive to different opinions. The candidates can freely join or depart from political coalitions, with the aim of maximizing their influence on the electorate.

Hedonic skill games also model the situation where some volunteers (agents) decide to join some charity organizations (coalitions) in order to contribute to some humanitarian activities (tasks). Those activities require some skills (e.g., medicine, logistics, translation, education, etc.) held by some agents. The volunteers are free to decide which organization to join, with the objective of being as helpful as possible.

1.1 Related Work

The game studied in this article falls into the family of *coalitional formation games* where the outcome is a coalition structure, i.e., a partition of the set of agents [Haj06]. Each agent has a preference relation over the coalition structures. In most cases, agents only care about their own coalition in the coalition structure. Such coalition formation games are called *hedonic* [DG80, AS16], a property shared by the game studied in this article. A generalization of hedonic games are *generalised group activity selection games* where we are given a set of activities and a set of agents. Each agent has to be assigned to one activity and agent’s preferences bear both on the activity she will be assigned to, and on the set of agents who will participate in the same activity. Specific scenarios of generalised group activity selection games have been studied in [BFF⁺19, DEK⁺18, IPE17].

Our work is mainly connected to the articles [BR08, BMJK10, BPR13]. Coalitional skill games (CSGs) were introduced in [BR08] (see also [BPR13]) and mainly studied from a cooperative point of view. The work in [BR08, BPR13] analyzes the expressiveness of CSGs and presents several restrictions that permit a concise yet good representation for these games. The authors also consider the computational complexity of several natural problems in WTSGs and other interesting subclasses.¹ Specifically, they study the complexity of the following issues: testing if an agent is a dummy or veto agent, computing the core and core-related solution concepts, and computing power indices such as the Shapley value and Banzhaf power index. The work in [BMJK10] considers the computational complexity of computing optimal coalition structures in CSGs (see also [AdK11] for this approach in numerous coalitional games including CSGs) and show that the problem is hard even for very restricted classes of games, i.e., Single Task Skill Games (STSGs) and WTSGs where tasks require at most two skills and skills are owned by at most two agents. However, they present a fixed parameter tractable algorithm for instances whose underlying structure has a bounded tree-width. Another positive result is the existence of a polynomial-time algorithm that computes a socially optimal coalition structure for WTSGs with a fixed number of skills [AdK11].

To the best of our knowledge, the strategic version of WTSGs that we consider in this paper was not considered before. Moreover, we emphasize that [BMJK10] mentioned in their conclusion that further game theoretic analysis is appropriate for the setting where agents are selfish and only care about their own utility.

A related model is the *Coalitional Resource Games* (CRGs) where agents wish to achieve various goals and are endowed with certain amounts of resources required to achieve these goals. CSGs and CRGs are related since performing a task in CSGs requires a coalition to have a certain set of skills, and achieving a goal in CRGs requires certain resources. However, there are important differences between CSGs and CRGs (see Section 5.4 of [BPR13] for a detailed comparison). [WD06] investigate the computational complexity of a number of natural decision problems for CRGs.

Another related topic is the *Task Oriented Team Formation problem*. We are given a task defined as a set of skills and agents possessing subsets of skills. For each pair of agents there is a cost that we have to pay if the two agents are selected in the team. The objective is to find a subset of agents that minimizes the overall cost subject to that every skill in the task is covered. To cover a skill, it is sufficient that one team member possesses that skill. [LLT09] show the hardness of this problem. Moreover, [OSC⁺15] consider the robust version of the problem and define a team k -robust (for a non-negative integer k) if it can be subject to the removal of any k agents and still complies with the skill requirements for the task.

Let us also mention the hedonic game of [SHB⁺11] where some wireless agents have to service a given set of entities by collecting and transmitting their data. Every such action is called a task. The agents form a coalition structure before executing the tasks. Though this model deals with tasks and coalition structures, it significantly differs from our problem because there is only one skill (data

¹Interesting subclasses of WTSGs that they consider are Task Count Skill Games (TCSGs) where each task has weight 1, and Single Task Skill Games (STSGs) where we have only a single task.

management), and each task is assigned to a single coalition.

Finally, a further related topic is the *hedonic expertise games* [ÇKO21] where we are given a global set of skills and a set of agents where each agent possesses a level of expertise in each of these skills. Agents form coalitions and agents belonging to the same coalition have the same utility which is defined as the sum of the maximum expertise for each skill among the agents of the coalition.

2 Model, Contribution, and Organization

The *hedonic skill game* is a strategic game composed of a set of agents $\mathcal{N} = \{1, \dots, n\}$, a set of tasks $\mathcal{T} = \{t_1, \dots, t_m\}$, and a set of skills $\mathcal{S} = \{s_1, \dots, s_k\}$. Each task t_j has a positive weight $w(t_j)$ and requires a non-empty set of skills $S(t_j) \subseteq \mathcal{S}$ in order to be executed. Every agent $\ell \in \mathcal{N}$ has a non-empty skill set $S(\ell) \subseteq \mathcal{S}$ and she decides which coalition she wants to be part of. In a state of the game σ , each coordinate σ_ℓ indicates the coalition that agent ℓ has decided to join. Every state σ induces a *coalition structure* (C_1, \dots, C_h) such that each C_i is the non-empty subset of agents who selected C_i . A coalition structure is a partition of \mathcal{N} , i.e., $\emptyset \neq C_i \subseteq \mathcal{N}$ for all i , $\cup_{i=1}^h C_i = \mathcal{N}$, and $C_i \cap C_{i'} = \emptyset$ for all (i, i') satisfying $1 \leq i < i' \leq h$. The number of coalitions of (C_1, \dots, C_h) is h . In this article, we interchangeably use the notion of state of the game and its induced coalition structure, as they both determine the necessary information about the agents' joint decisions.

Every coalition C_i holds the set of skills $S(C_i) := \cup_{\ell \in C_i} S(\ell)$, which is the union of the skills of its members. A coalition C_i is able to perform a task t_j if its members have all the needed skills, i.e., when $S(t_j) \subseteq S(C_i)$. Let $T(C_i)$ denote the subset of tasks that C_i can perform.

$$T(C_i) := \{t \in \mathcal{T} \mid S(t) \subseteq S(C_i)\}$$

A task t can be performed by more than one coalition if several coalitions contain agents who possess the skills that t requires. However, a coalition executes a task at most once.

Depending on the situation that the game models, the number of coalitions of any coalition structure can be either unconstrained, or upper bounded by some given parameter that we denote by q . Since the number of coalitions of a coalition structure cannot exceed the number of agents, $q \leq n$ can be assumed without loss of generality.

Let us now define the utility of the agents. The hedonic nature of our skill game comes from the fact that the utility of every agent only depends on the agents that are in her coalition.

The weight $w(t)$ of every task t that a coalition C can perform is distributed among C 's members having at least one skill within $S(t)$. Concretely, for all skill $s \in S(t)$, every agent $\ell \in C$ such that $s \in S(\ell)$ receives a reward of $\frac{w(t)}{|S(t)| \cdot |\{a \in C \mid s \in S(a)\}|}$. The total sum of the rewards that an agent receives constitutes her utility. Thus, the utility of agent $\ell \in C$ is equal to

$$\sum_{t \in T(C)} \sum_{s \in S(t) \cap S(\ell)} \frac{w(t)}{|S(t)| \cdot |\{a \in C \mid s \in S(a)\}|}.$$

Given this definition, the sum of the agents' utilities within a coalition C is always equal to the total weight of the tasks performed by C . Moreover, the sum over \mathcal{N} of the agents' utilities is equal to the total weight of the tasks performed in the whole coalition structure.² As a notation, u_ℓ denotes agent ℓ 's utility function.

The social welfare associated with a state σ , denoted by $SW(\sigma)$, is defined as the sum of the agents' utilities (a.k.a. utilitarian social welfare). From the above discussion, $SW(\sigma)$ is also equal to the total weight of all the executed tasks. Thus, if σ induces the coalition structure (C_1, \dots, C_h) , then it holds that

$$SW(\sigma) = \sum_{\ell \in \mathcal{N}} u_\ell(\sigma) = \sum_{i=1}^h \sum_{t \in T(C_i)} w(t). \quad (1)$$

A social optimum is a state σ^* for which $SW(\sigma^*)$ is maximum.

Example 1. Suppose there are three skills $\{s_1, s_2, s_3\}$, two tasks $\{t_1, t_2\}$ of weights $w(t_1) = 6$ and $w(t_2) = 4$ requiring skills $S(t_1) = \{s_1, s_2\}$ and $S(t_2) = \{s_2, s_3\}$, and three agents having skills $S(1) =$

²The weight of a task is counted x times if it is performed by x distinct coalitions.

$\{s_1, s_2\}$, $S(2) = \{s_1, s_3\}$, and $S(3) = \{s_1\}$. If the three agents are in the same coalition, then both tasks can be performed: agent 1 has utility 6, agent 2 has utility 3, agent 3 has utility 1, and the social welfare is 10. If agent 2 is alone and agents 1 and 3 are together in the same coalition, then only the task t_1 can be performed by the coalition with agents 1 and 3: agent 1 has utility $9/2$, agent 2 has utility 0, agent 3 has utility $3/2$, and the social welfare is 6.

A (pure) *Nash stable* outcome (a.k.a. Nash equilibrium in pure strategy) is a situation where no agent can deviate from her current coalition and strictly increase her utility (i.e., no agent possesses a better response).

Example 2. Suppose there are two skills $\{s_1, s_2\}$, a single task t_1 of weight $w(t_1) = 2$ requiring skills $S(t_1) = \{s_1, s_2\}$, and two agents having skills $S(1) = \{s_1, s_2\}$ and $S(2) = \{s_2\}$. If the two agents are in the same coalition, then the task t_1 can be performed: agent 1 has utility $3/2$ whereas agent 2 has utility $1/2$. If agent 2 is alone, then the task cannot be performed in her coalition, and agent 2's utility is 0. If agent 1 is alone, then the task can be performed in her coalition, and agent 1's utility is 2.

Example 2 disproves the existence of a Nash stable outcome, even if the number of possible coalitions is limited to two. Indeed, agent 1 prefers being alone over being with agent 2, and agent 2 prefers being with agent 1 over being alone.

Observe that in Example 2, agent 1 and the task both have more than one skill. It is interesting to consider specific instances of hedonic skill games for which it is possible to prove that Nash equilibria are guaranteed to exist. An agent (resp., a task) is said to be *singleton* if she has (resp., it requires) a single skill. We will differentiate the instances solely composed of singleton agents (resp., singleton tasks), and the others called *general agents instances* (resp., *general tasks instances*) where the agents (resp., tasks) are not necessarily singleton. As we will see later, this distinction has a significant influence on the various aspects of the hedonic skill game: existence of a Nash stable outcome, convergence of the better response dynamics, complexity of computing socially optimal coalition structures, and the price of anarchy. In addition, the upper bound q on the number of possible coalitions of the coalition structure can also play a role. Since studying all the aforementioned aspects for every possible value of q represents a huge amount of work, the focus is put on the two extreme cases, namely $q = n$ (the number of coalitions of the coalition structure is unconstrained) and $q = 2$.

2.1 Contribution

We first show that for general instances such that $q \geq 3$, the problem of deciding whether an instance of the hedonic skill game admits a Nash stable outcome is NP-complete (Theorem 1). We then provide a complete picture about the existence of Nash stable outcomes in the hedonic skill game. Specifically, we demonstrate that a Nash stable outcome always exists if all the agents are singleton (Theorem 3 for $q = n$ and Proposition 2 for $q = 2$) or if all the tasks are singleton (Corollary 1). These results are obtained with different techniques. We notice that Example 2 covers the remaining cases (namely, non-existence of a Nash stable outcome for general agents/tasks instances) for all $q \geq 2$. We also show that computing a Nash stable outcome is a **PLS**-complete problem (Theorem 5) with singleton tasks, indicating that it is a difficult task, but the problem is in **P** if the agents are singleton (Theorem 3 for $q = n$ and Proposition 2 for $q = 2$).

Moreover, as in previous works on CSGs, we adopt the utilitarian social welfare, which is defined as the sum of the agents' utilities, as measure of the well-being of the system. It is known that computing a socially optimal coalition structure is **NP**-hard, even in restricted cases [BMJK10], i.e., for Single Task Skill Games (STSGs) where there is only one task and for WTSGs where tasks require at most two skills and skills are owned by at most two agents. In this work we show that a social optimum can be built in polynomial time for singleton agents instances when $q = n$ (Theorem 2) and when $q = 2$ (Proposition 2). Concerning singleton tasks instances, we prove that maximizing the social welfare is **NP**-hard when $q = 2$ and polynomially time solvable when $q = n$ (Theorem 6).

We further study agents dynamics. Specifically, starting from any coalition structure, the better response dynamics (BRD, in short) consists of repeatedly applying better responses by single agents³ as long as it is possible. If BRD eventually stops, then we say that it converges and the final state must be a Nash stable outcome. Otherwise, we say that the dynamics cycles. We prove that BRD

³Choose one deviation arbitrarily if there is more than one.

PoA	singleton agents	singleton tasks
$q = 2$	PoA=3 if $\tau = 2$, PoA=5 if $\tau = 3$, and unbounded PoA if $\tau > 3$ (Theorem 8)	PoA=4/3 (Theorem 9)
$q = n$	$\frac{n}{2} \leq \text{PoA} \leq 2n$ if $\tau = 2$, and unbounded PoA if $\tau > 2$ (Theorem 7)	PoA=1

Table 1: Price of Anarchy (PoA for short).

always converges if the tasks are singleton (Theorem 4) but it can cycle with general tasks, even if agents are singleton (Proposition 1).

Finally, we study the quality of Nash stable outcomes by resorting to the price of anarchy (PoA, in short), which is the largest value, over all instances, of the social welfare (SW) of a Nash stable outcome divided by the optimal social welfare [KP99]. Our results, summarized in Table 1, show that the PoA is sometimes sensitive to a parameter τ defined as the maximum number of skills required by a task.

2.2 Organization of this Article

We first study the complexity of deciding whether a given general instance possesses a Nash stable outcome or not (Section 3). We then consider the case of singleton agents (Section 4), followed by the case of singleton tasks (Section 5) where we study the existence and computation of Nash stable outcomes when $q \in \{2, n\}$, and the convergence of BRD. Section 6 is dedicated to the analysis of the price of anarchy. We conclude and provide a list of possible future works in Section 7.

3 Hardness of deciding if a Nash stable outcome exists

We know from Example 2 that some instances fail to have a Nash stable outcome. Therefore, it is natural to settle the complexity of deciding whether a given general instance possesses a Nash stable outcome or not. In the following theorem we show that deciding whether an instance of the hedonic skill game admits a Nash stable outcome is a hard problem when $q \geq 3$.

Theorem 1. *Deciding whether an instance of the hedonic skill game admits a Nash stable outcome is an NP-complete problem when $q \geq 3$.*

Sketched proof. The starting point is the NP-complete PARTITION problem. Given n positive integers a_1, \dots, a_n such that $\sum_{i=1}^n a_i = 2B$, decide if the numbers can be partitioned in two subsets which both sum up to B .

Given an instance \mathcal{I}_P of PARTITION, an instance \mathcal{I}_H of the hedonic skill game is built as follows.

- create an agent i with skills $\{s_i, s'_i, s''_i\}$ for all $i \in [n]$,
- create two agents with skills $\{s_0, s_1, s_2, \dots, s_n\}$ and call these agents $n+1$ and $n+2$, respectively,
- create $q-2$ agents with skills $\{s_0, s_1, s_2, \dots, s_n, s_{n+1}\} \cup \{s''_1, s''_2, \dots, s''_n\}$ and call these agents $n+3, \dots, n+q$,
- create a task of weight $4a_i$ requiring skills $\{s_i, s'_i\}$ for all $i \in [n]$,
- create a task of weight $\varepsilon = 0.4$ requiring skill s''_i for all $i \in [n]$,
- create a task of weight $4B - \varepsilon$ requiring skills $\{s_0, s_{n+1}\}$.

We claim that \mathcal{I}_H admits a Nash stable outcome if and only if \mathcal{I}_P is a yes instance. Indeed, there is a correspondence between Nash stable outcomes and subsets of indices X such that $\sum_{i \in X} a_i = \sum_{i \in [n] \setminus X} a_i = B$. The shape of a Nash stable outcome and its corresponding subset X is as follows. Agent i such that $i \in X$ is in the first coalition, agent i such that $i \in [n] \setminus X$ is in the second coalition,

agent $n + 1$ is in the first coalition, agent $n + 2$ is in the second coalition, and every remaining coalition from 3 to q admits exactly one agent amongst $\{n + 3, \dots, n + q\}$. \square

Despite that the problem is NP-complete in general, in the remainder of the paper we propose a characterization of the instances admitting a Nash stable outcome. We conclude this section by noticing that it remains open to study the complexity of deciding whether a given general instance possesses a Nash stable outcome or not for the case of $q = 2$.

4 Nash Stability with Singleton Agents

In this section we suppose that every agent has a single skill. The tasks can require more than one skill.

4.1 Unconstrained Number of Coalitions

We first consider the case where $q = n$. We start by showing in Theorem 2 that it is possible to compute in polynomial time a coalition structure which maximizes the social welfare (this problem is known to be NP-hard in general [BMJK10]). We build upon this first result to show in Theorem 3 how to compute a Nash stable outcome in polynomial time.

Theorem 2. *Finding a coalition structure which maximizes the social welfare SW can be done in polynomial time for singleton agents instances when $q = n$.*

Proof. The coalition structure is built from scratch with the help of a set of agents X . At the beginning X is equal to \mathcal{N} . Let $i = 1$. Afterwards, do the following steps repeatedly until $X = \emptyset$: create a coalition $C_i \subseteq X$ such that $|C_i| = |S(X)|$ and $S(C_i) = S(X)$, remove C_i from X , and $i \leftarrow i + 1$.

At each step, the algorithm creates a new coalition which has one representative (chosen arbitrarily) for every skill present in X . Let (C_1, \dots, C_t) denote the coalition structure built by the algorithm. We clearly have $S(C_1) \supseteq S(C_2) \supseteq \dots \supseteq S(C_t)$.

For every skill $s \in \mathcal{S}$, let $n(s)$ be the total number of agents having skill s in the instance. For every subset of skills $Y \subseteq \mathcal{S}$, let $\mu(Y)$ denote $\min_{s \in Y} n(s)$. Each task $t \in T$ can be performed at most $\mu(S(t))$ times because at most $\mu(S(t))$ coalitions can contain all the needed skills.

By construction of (C_1, \dots, C_t) , every skill s appears in coalitions $C_1, \dots, C_{n(s)}$. Thus, for every non-empty $Y \subseteq \mathcal{S}$, we have $Y \subseteq C_i$ for $i = 1.. \mu(Y)$. It follows that (C_1, \dots, C_t) maximizes the social welfare because every task t is executed its maximum number of times $\mu(S(t))$. \square

Note that Theorem 2 can be generalized to any value of q .

Theorem 3. *A Nash stable outcome always exists and it can be computed in polynomial time for singleton agents instances.*

Proof. We consider a dedicated dynamics D which starts from a particular coalition structure built in polynomial time, and at each step, D selects a specific deviation. We will show that after a polynomial number of deviations, a Nash stable outcome is reached.

For every possible state, the corresponding coalition structure is represented with a $k \times n$ table Y where each row corresponds to a skill, each column is associated with a coalition, and the entry $y_{i,j}$ of Y is equal to the number of agents having skill s_i and belonging to coalition C_j . Given Y , we can define another $k \times n$ table L whose entry $L_{i,j}$ is the sum of rewards shared by the (possibly prospective) agents of skill s_i in coalition C_j , given that the skills $s_{i'} \neq s_i$ already present in C_j are in accordance with Y . More precisely, we consider that every skill $s_{i'}$ such that $i' \neq i$ is present in C_j if, and only if, $y_{i',j} > 0$. Then, $L_{i,j}$ is the sum of rewards that agents of skill s_i would collectively receive if they were present in C_j . Therefore, no matter what $y_{i,j}$ currently is, if exactly $x > 0$ agents having skill s_i are in C_j , then each of them would have utility $L_{i,j}/x$. In the following, if an agent belongs to coalition C_j and performs an improving move by joining coalition $C_{j'}$, then we call C_j and $C_{j'}$ the *departing* and *arrival* coalitions, respectively. A single agent move (simply called “move” afterwards) from C_j to $C_{j'}$ is *left* if $j > j'$ (resp., *right* if $j < j'$). A move is *best* (resp., *better*) if it is a best (resp., better) response. Given a and b satisfying $1 \leq a < b \leq n$, $C_a \xleftarrow{\text{best}} C_b$ denotes a best left move by an agent having skill s_i from departing coalition C_b to arrival coalition C_a .

The initial state of the dynamics D is the coalition structure built in the proof of Theorem 2. Thus, each row of Y initially consists of consecutive 1s, followed by consecutive 0s. While the current state is not a Nash stable outcome, the next move of the dynamics D must be a best left move invariably selected as follows.

1. For every skill index i such that at least one agent having skill s_i wants to make a best left move, do: Each time there are two possible moves $C_a \xleftarrow{\text{best } i} C_b$ and $C_{a'} \xleftarrow{\text{best } i} C_b$ such that $a < a' < b$, discard $C_{a'} \xleftarrow{\text{best } i} C_b$. Each time there are two possible moves $C_a \xleftarrow{\text{best } i} C_{b'}$ and $C_a \xleftarrow{\text{best } i} C_b$ such that $a < b' < b$ and $y_{i,b'} = y_{i,b}$, discard $C_a \xleftarrow{\text{best } i} C_{b'}$.
2. Within the best left moves that survived step 1, execute the one whose departing coalition has smallest index. Break ties arbitrarily.

The dynamics D that we are going to describe maintains the following invariants:

$$y_{i,j} \geq y_{i,j+1} \text{ for all pair } (i,j) \quad (2)$$

$$L_{i,j} \geq L_{i,j+1} \text{ for all pair } (i,j) \quad (3)$$

$$\text{No better right move exists} \quad (4)$$

Invariants (2) and (3) are trivially satisfied by the initial state. The remainder of the proof relies on three intermediate Lemmas.

Lemma 1. *The initial state of D satisfies Invariant (4).*

It follows that if the initial state is not a Nash stable outcome, then the first move of D is a best left move. For the guaranteed existence of a Nash stable outcome, it remains to prove that the three invariants are maintained throughout the execution of D . Indeed, either a Nash stable outcome is reached, or the previous move (which is a best left move) never leaves the possibility for an agent to move to the right in the table Y , meaning that D eventually stops.

Lemma 2. *The dynamics D maintains Invariants (2) and (3).*

Lemma 3. *A best left move does not trigger the existence of a better right move, i.e., Invariant (4) is maintained.*

The convergence occurs after a polynomial number of steps because D solely performs left moves (there are n agents and each one can move at most $n - 1$ times), and determining the next move performed by D (or that a Nash stable outcome is reached) can be done in $O(kn^2)$ by checking, for every skill index i , if it is worth moving from C_b to C_a , where $1 \leq a < b \leq n$. \square

In the proof of Theorem 3, the existence of a Nash stable outcome is obtained with a specific dynamics applied on a particular initial coalition structure. However, it is legitimate to ask whether a Nash stable outcome can be reached by a natural dynamics such as the better response dynamics (BRD).

Proposition 1. *There exists an instance of the hedonic skill games with singleton agents where BRD cycles.*

Proof. Let us present an instance and a cyclic sequence of better response deviations. The instance consists of 8 agents and 7 tasks. There are 3 agents with skill s_1 , and one agent per skill s_i such that $2 \leq i \leq 6$. The following list gives the details of the 7 tasks, i.e., the name of the task, its weight, and the set of skills that it requires: $(t_{12}, 78, \{s_1, s_2\})$, $(t_{14}, 96, \{s_1, s_4\})$, $(t_{23}, 114, \{s_2, s_3\})$, $(t_{15}, 102, \{s_1, s_5\})$, $(t_{16}, 54, \{s_1, s_6\})$, $(t_{123}, 18, \{s_1, s_2, s_3\})$, $(t_{134}, 180, \{s_1, s_3, s_4\})$.

The 8 states of the cycle are reported in Table 2, where each multiset represents a coalition and the skills of its members. The first deviation is done from C_1 to C_2 by an agent having skill s_1 . The initial and final utilities of the agent are $78/2 + 18/3 = 45$ and $96/2 = 48$, respectively. The second deviation is done from C_1 to C_2 by an agent having skill s_3 . The initial and final utilities of the agent are $114/2 = 57$ and $180/3 = 60$, respectively. The third deviation is done from C_3 to C_2 by an agent having skill s_1 . The initial and final utilities of the agent are $102/2 = 51$ and $96/4 + 180/6 = 54$, respectively. The fourth deviation is done from C_4 to C_2 by an agent having skill s_1 . The initial and final utilities of the agent are $54/2 = 27$ and $96/6 + 180/9 = 36$, respectively. The fifth deviation

Table 2: The BRD cycle of Proposition 1

C_1	C_2	C_3	C_4
$\{\{s_1, s_2, s_3\}\}$	$\{\{s_4\}\}$	$\{\{s_1, s_5\}\}$	$\{\{s_1, s_6\}\}$
$\{\{s_2, s_3\}\}$	$\{\{s_1, s_4\}\}$	$\{\{s_1, s_5\}\}$	$\{\{s_1, s_6\}\}$
$\{\{s_2\}\}$	$\{\{s_1, s_3, s_4\}\}$	$\{\{s_1, s_5\}\}$	$\{\{s_1, s_6\}\}$
$\{\{s_2\}\}$	$\{\{s_1, s_1, s_3, s_4\}\}$	$\{\{s_5\}\}$	$\{\{s_1, s_6\}\}$
$\{\{s_2\}\}$	$\{\{s_1, s_1, s_1, s_3, s_4\}\}$	$\{\{s_5\}\}$	$\{\{s_6\}\}$
$\{\{s_1, s_2\}\}$	$\{\{s_1, s_1, s_3, s_4\}\}$	$\{\{s_5\}\}$	$\{\{s_6\}\}$
$\{\{s_1, s_2, s_3\}\}$	$\{\{s_1, s_1, s_4\}\}$	$\{\{s_5\}\}$	$\{\{s_6\}\}$
$\{\{s_1, s_2, s_3\}\}$	$\{\{s_1, s_4\}\}$	$\{\{s_5\}\}$	$\{\{s_1, s_6\}\}$

is done from C_2 to C_1 by an agent having skill s_1 . The initial and final utilities of the agent are $96/6 + 180/9 = 36$ and $78/2 = 39$, respectively. The sixth deviation is done from C_2 to C_1 by an agent having skill s_3 . The initial and final utilities of the agent are $180/3 = 60$ and $18/3 + 114/2 = 63$, respectively. The seventh deviation is done from C_2 to C_4 by an agent having skill s_1 . The initial and final utilities of the agent are $96/4 = 24$ and $54/2 = 27$, respectively. The last deviation goes back to the initial state. It is done from C_2 to C_3 by an agent having skill s_1 . The initial and final utilities of the agent are $96/2 = 48$ and $102/2 = 51$, respectively. \square

Nevertheless, BRD always converges in singleton agents instances where every skill is possessed by a single agent when $q = n$. This is due to the existence of a potential.

4.2 At Most Two Coalitions

Theorem 3 holds for an unconstrained number of coalitions. The technique consists of applying specific dynamics on a particular initial coalition structure. However, when q is restricted, the technique is not guaranteed to work if the aforementioned initial state requires more than q coalitions. Therefore, we show in the following how a Nash stable state can be efficiently computed when $q = 2$.

Proposition 2. *When $q = 2$, hedonic skill games with singleton agents always have a pure Nash stable outcome which can be computed in polynomial time. Moreover, computing a state that maximizes the social welfare in this case can be done in polynomial time.*

Sketched proof. For every skill, balance the corresponding agents on the two coalitions to obtain a state that maximizes the social welfare. The first coalition always gets more agents of a given skill if their number is odd. Afterwards, possible profitable deviations are always from the second coalition towards the first coalition. \square

The question whether BRD converges when $q = 2$ (and agents are singleton) is an intriguing open problem. Note that the instance given in Proposition 1 contains 4 coalitions, so the proposition holds when $q \geq 4$.

5 Nash Stability with Singleton Tasks

In this section we suppose that every task requires exactly one skill. There is a one-to-one correspondence between skills and tasks: each skill s is associated with a task t_s of weight w_s requiring s . The agents can have more than one skill.

Unlike the previous section, we do not approach the existence of a Nash stable outcome starting from a social optimum. We show that BRD always converges, proving that a Nash stable outcome always exists. However, we demonstrate that it is unlikely that a Nash stable outcome can be computed efficiently.

Theorem 4. *For every parameter q , BRD always converges for hedonic skill games with singleton tasks.*

The proof of Theorem 4 consists of showing that every instance is a *congestion game* [MS96]. Congestion games are central to the field of algorithmic game theory [NRTV07, vS21]. They always admit an *exact potential function* (namely, Rosenthal’s potential function [Ros73]). An exact potential function Φ associates a real value with every state. Each time an agent ℓ makes a profitable deviation that turns the current state σ into a new state σ' , it holds that $\Phi(\sigma') - \Phi(\sigma) = u_\ell(\sigma') - u_\ell(\sigma)$. Every local optimum of Φ , defined as a state such that no single agent can profitably deviate, is Nash stable. Thus, BRD is known to always converge in congestion games.

A congestion game consists of a set of agents, a set of resources R , and a function $g : R \times \mathbb{N} \rightarrow \mathbb{R}$. The strategy space of every agent is a collection of subsets of R . If exactly x agents have a given resource r in their strategy, then the gain associated with r for every user of r is $g(r, x)$. The payoff of an agent is the total sum of her gains, over the subset of resources contained in her strategy. Rosenthal’s potential of σ is defined as $\Phi(\sigma) = \sum_{r \in R} \sum_{x=1}^{n_r(\sigma)} g(r, x)$, where $n_r(\sigma)$ is the number of agents having resource r in their strategies.

Proof of Theorem 4. Start from an instance of the hedonic skill game and create a set R of $q|\mathcal{S}|$ resources as follows: r_i^j with $i \in \{1, \dots, |\mathcal{S}|\}$ and $j \in \{1, \dots, q\}$. For every agent $\ell \in \mathcal{N}$, let $I(\ell)$ denote the indices of agent ℓ ’s skills, i.e., $I(\ell) := \{i \mid s_i \in S(\ell)\}$. The strategy space of every agent ℓ includes $\{r_i^j \mid i \in I(\ell)\}$ for every $j \in \{1, \dots, q\}$. The function $g(r_i^j, x)$ is defined as w_i/x where w_i is the weight of the task requiring only skill s_i .

If an agent ℓ plays $\{r_i^j \mid i \in I(\ell)\}$ for some $j \in \{1, \dots, q\}$, then this means that ℓ is in coalition C_j . If ℓ has skill s_i , and in total x agents having skill s_i belong to coalition C_j , then each of these agents gets w_i/x , which corresponds to the reward associated with the task solely requiring skill s_i . \square

The following corollary is a direct consequence of Theorem 4.

Corollary 1. *Hedonic skill games with singleton tasks always admit a Nash stable outcome.*

One can directly use BRD for computing a Nash stable outcome, but the number of steps before convergence is not guaranteed to be polynomial in the parameters of the game. A different algorithm may be used for efficiently computing a Nash stable outcome but we show below, with tools from local search theory, that its existence is unlikely when $q = 2$.

Theorem 5. *Computing a Nash stable outcome in hedonic skill games with singleton tasks is a PLS-complete problem when $q = 2$.*

Proof. MAX CUT FLIP is the problem of computing a solution of MAX CUT⁴ such that flipping a vertex (i.e., moving it to the other part of the bi-partition) does not increase the total weight of the cut. Such a stable solution is called a local optimum and MAX CUT FLIP is known to be PLS-complete [SY91].

The problem of computing a Nash stable outcome of the hedonic skill games with singleton tasks and $q = 2$ is clearly in PLS since one can easily build an initial solution (e.g., the grand coalition), and one can decide in polynomial time if it is profitable for an agent to deviate to the other coalition.

Let us show that MAX CUT FLIP is PLS-reducible to the computation of a Nash stable outcome of the hedonic skill game with singleton tasks and $q = 2$.

Every instance \mathcal{I} of MAX CUT FLIP is mapped to an instance $f(\mathcal{I})$ of hedonic skill games with $q = 2$ as follows. Each vertex v_i is associated with an agent i . Each edge (v_i, v_j) corresponds to a skill s_{ij} and that skill is solely held by agents i and j . Each edge (v_i, v_j) also corresponds to a task requiring skill s_{ij} and the weight of the task is equal to the weight of the edge. Each part of the bi-partition corresponds to a coalition so that every state of $f(\mathcal{I})$ is naturally associated with a solution of \mathcal{I} . It remains to prove that every Nash stable outcome of $f(\mathcal{I})$ is a local optimum of \mathcal{I} .

In a Nash stable outcome, no agent can profitably deviate. The utility of agent i is equal to $\frac{w_{in}(i)}{2} + w_{out}(i)$ where $w_{in}(i)$ denotes the weight of the tasks associated with edges (v_i, v_j) such that agent j is in the same coalition as i , while $w_{out}(i)$ stands for the weight of the tasks associated with edges (v_i, v_k) such that agent k is not in the same coalition as i . If agent i deviates, then her new utility becomes $w_{in}(i) + \frac{w_{out}(i)}{2}$. Thus, every agent i satisfies $w_{out}(i) \geq w_{in}(i)$ in a Nash stable outcome.

The weight of the cut is $w_{out}(i) + w_{rest}(i)$ where $w_{rest}(i)$ denotes the total weight of the edges (v_j, v_k) such that $i \notin \{j, k\}$, and j and k are in distinct coalitions. If i deviates, then the weight of the

⁴Given a simple graph $G = (V, E)$ with positive weights on its edges, find a bi-partition (V_1, V_2) of V such that the total weight of the edges having an endpoint in V_1 and the other endpoint in V_2 , a.k.a. *the cut*, is maximized [Kar72].

cut becomes $w_{in}(i) + w_{rest}(i)$. Since $w_{out}(i) \geq w_{in}(i)$ holds for all i , we get that $w_{out}(i) + w_{rest}(i) \geq w_{in}(i) + w_{rest}(i)$. In other words, no vertex flipping increases the weight of the cut. \square

Observe that building a Nash stable outcome is an easy task when the number of coalitions is unconstrained ($q = n$): if every agent is alone in her coalition, then nobody wants to deviate since everybody enjoys her maximum possible utility. Let us conclude this section with computational complexity of optimizing the social welfare of the hedonic skill game with singleton tasks.

Theorem 6. *Regarding the hedonic skill game with singleton tasks, maximizing the social welfare SW is an NP-hard problem when $q = 2$, and polynomial time solvable when $q = n$*

6 Price of Anarchy

This section is devoted to the *price of anarchy* (PoA) of hedonic skill games. The PoA is the largest value (over the entire set of instances of the game) taken by the maximum social welfare divided by the social welfare of the worst Nash stable state. We consider the PoA with respect to the utilitarian social welfare SW which is defined as the sum of the agents' utilities (cf. (1) in Section 2). Instances of hedonic skill games having non-singleton agents and non-singleton tasks at the same time are not considered because the existence of a Nash stable outcome is not guaranteed (cf. Example 2). Our results focus on the two extreme values of q , namely 2 and n .

When both agents and tasks are singleton, then the PoA is 1 (Proposition 3). The PoA of the remaining singleton agents instances depends on the maximum number of skills required by a task (denoted by τ)⁵ and is as follows: $O(n)$ when $\tau = 2$ and $q = n$ (Theorem 7), $O(1)$ when $\tau \in \{2, 3\}$ and $q = 2$ (Theorem 8), unbounded otherwise (Theorems 7 & 8). Concerning the singleton tasks instances, we show that the PoA is $4/3$ when $q = 2$ (Theorem 9), and 1 when $q = n$. All the results are tight except for the case of singleton agents, $\tau = 2$ and $q = n$ where the result is asymptotically tight.

Proposition 3. *For the case of singleton tasks and singleton agents, the PoA of hedonic skill games is 1 for all q .*

The omitted proof of Proposition 3 shows that the singleton task associated with any given skill is executed its maximum possible number of times in any Nash stable outcome.

6.1 Singleton Agents

In this section we consider singleton agents. We first consider the case of $q = n$ (Theorem 7) and then the case of $q = 2$ (Theorem 8). For both cases, we provide a tight or an asymptotically tight analysis on the PoA.

Theorem 7. *If $q = n$, then the PoA of hedonic skill games with singleton agents is $\Theta(n)$ when $\tau = 2$, and unbounded when $\tau > 2$.*

Proof. Let us begin with the case $\tau = 2$.

Fix any Nash stable outcome σ . We say that a task t is *active* in σ if there exists at least one coalition $C \in \sigma$ which is able to perform the task t , i.e., $S(t) \subseteq S(C)$, and *non-active* otherwise. Let us denote the set of active tasks in σ by $A(\sigma)$, and the set of non-active tasks in σ by $\overline{A(\sigma)} = \mathcal{T} \setminus A(\sigma)$. Note that in any Nash stable outcome, all the singleton tasks are active, while any task of size two (i.e., requiring two skills) can be either active or non-active. We get that

$$SW(\sigma) = \sum_{\ell \in \mathcal{N}} u_{\ell}(\sigma) \geq \sum_{t: t \in A(\sigma) \wedge |S(t)|=1} w(t) + \sum_{t: t \in A(\sigma) \wedge |S(t)|=2} w(t). \quad (5)$$

For any agent $\ell \in \mathcal{N}$, let $T_{\ell}(\sigma) = \{t \in \overline{A(\sigma)} : S(\ell) \subset S(t) \wedge |S(t)| = 2\}$ be the set of non-active tasks of σ of size two that require skill $S(\ell)$. Moreover, let $m_{\ell}(\sigma) = \arg \max_{t \in T_{\ell}(\sigma)} w(t)$ be the task belonging to $T_{\ell}(\sigma)$ of maximum weight. When $T_{\ell}(\sigma)$ is empty, we consider that $m_{\ell}(\sigma)$ is a dummy task of weight zero. It is easy to see that, for any agent $\ell \in \mathcal{N}$, it holds that $u_{\ell}(\sigma) \geq \frac{w(m_{\ell}(\sigma))}{2}$. Indeed, if $u_{\ell}(\sigma) < \frac{w(m_{\ell}(\sigma))}{2}$, then σ is not a Nash stable outcome since agent ℓ can perform an improving move

⁵Hence, $\tau = 1$ corresponds to singleton tasks.

(assuming that $m_\ell(\sigma)$ requires skills $\{S(\ell), S(a)\}$, for some agent a , the improving move of agent ℓ is to join any coalition having an agent with skill $S(a)$ by making the task $m_\ell(\sigma)$ active and, in this way, getting utility of at least $\frac{w(m_\ell(\sigma))}{2}$). Therefore, we get that

$$SW(\sigma) = \sum_{\ell \in \mathcal{N}} u_\ell(\sigma) \geq \sum_{\ell \in \mathcal{N}} \frac{w(m_\ell(\sigma))}{2}. \quad (6)$$

In an optimal coalition structure, any singleton task can be performed by at most n coalitions, and any active task of size two can be performed by at most $\frac{n}{2}$ coalitions. Moreover, let us denote by $SW(\overline{A(\sigma)})$ the social welfare that any optimal solution extracts from non-active tasks. We have

$$SW(\overline{A(\sigma)}) \leq \frac{1}{2} \sum_{\ell \in \mathcal{N}} \sum_{t \in T_\ell(\sigma)} w(t) \leq \frac{1}{2}(n-1) \sum_{\ell \in \mathcal{N}} w(m_\ell(\sigma)), \quad (7)$$

where the first inequality holds since each non-active task of size two is counted at least twice, and the second inequality holds since, for any agent $\ell \in \mathcal{N}$, there exist in σ at most $n-1$ non-active tasks of size two that require skill $S(\ell)$ (because there are at most n coalitions) and each of them has weight at most $w(m_\ell(\sigma))$.

Therefore, given any optimal coalition structure σ^* , we have that

$$\begin{aligned} SW(\sigma^*) &= \sum_{\ell \in \mathcal{N}} u_\ell(\sigma^*) \leq (n \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=1} w(t)) + \left(\frac{n}{2} \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=2} w(t)\right) + SW(\overline{A(\sigma^*)}) \leq \\ &\leq (n \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=1} w(t)) + \left(\frac{n}{2} \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=2} w(t)\right) + \left(\frac{1}{2}(n-1) \sum_{\ell \in \mathcal{N}} w(m_\ell(\sigma^*))\right) \end{aligned} \quad (8)$$

By combining inequalities (5), (6) and (8) we get

$$\frac{SW(\sigma^*)}{2SW(\sigma)} \leq \frac{(n \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=1} w(t)) + (\frac{n}{2} \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=2} w(t)) + (\frac{1}{2}(n-1) \sum_{\ell \in \mathcal{N}} w(m_\ell(\sigma^*)))}{\sum_{t: t \in A(\sigma^*) \wedge |S(t)|=1} w(t) + \sum_{t: t \in A(\sigma^*) \wedge |S(t)|=2} w(t) + \sum_{\ell \in \mathcal{N}} \frac{w(m_\ell(\sigma^*))}{2}} \leq n,$$

which implies that

$$\frac{SW(\sigma^*)}{SW(\sigma)} \leq 2n. \quad (9)$$

Since inequality (9) holds for any Nash stable outcome, and in particular for the one with worse (i.e., smallest) social welfare, the upper bound on the PoA follows.

For the lower bound ($\tau = 2$), suppose the number of agents n is even. There are $n/2$ agents with skill s_1 and $n/2$ agents with skill s_2 . There is a single task t of weight 1 which requires both s_1 and s_2 . The grand coalition is a Nash stable outcome in which every agent has utility $1/n$. If an agent deviates, then she is alone in a new coalition, and her utility is 0. The social welfare of the grand coalition is 1. Consider the coalition structure composed of $n/2$ coalitions, each one having an agent with s_1 and an agent with s_2 . The social welfare of this state is $n/2$. Therefore, $\text{PoA} \geq \frac{n}{2}$.

Finally, consider the case $\tau > 2$. Suppose n is a multiple of τ . There are τ skills, and n/τ agents per skill. There is a single task of weight 1 requiring all the skills. The state where all the agents are isolated (i.e., they form singleton coalitions) is a Nash stable outcome of null social welfare. The optimal social welfare is reached when n/τ coalitions of size τ are formed (each such coalition contains all the skills). Thus, the PoA is unbounded. \square

Theorem 8. *If $q = 2$, then the PoA of hedonic skill games with singleton agents is 3 when $\tau = 2$, $\text{PoA} = 5$ when $\tau = 3$, and the PoA is unbounded when $\tau > 3$.*

6.2 Singleton Tasks

In this section we consider singleton tasks. We first notice that when the number of coalitions is unconstrained ($q = n$), the PoA of hedonic skill games with singleton tasks is trivially equal to 1. This is because every agent can reach her maximum possible utility if she chooses to be alone in a coalition, and such a state is Nash stable. In the following theorem we deal with the case of $q = 2$ and provide a tight analysis on the PoA.

Theorem 9. *The PoA of hedonic skill games with singleton tasks is $4/3$ when $q = 2$.*

Proof. Fix an instance and consider a Nash stable outcome σ and a socially optimal outcome σ^* . Suppose there is a skill s that only one agent, say i , owns. The instance can be modified by deleting skill s because the PoA can only increase, and σ remains a Nash stable outcome (only the utility of agent i is decreased by the weight of the task associated with s , for every strategy that i can take). Therefore, we can suppose w.l.o.g. that no skill is owned by a single agent.

Let W_1 (resp., W_2) be the total weight of the tasks executed once (resp., twice) under σ . Let us prove the following intermediate result.

$$W_2 \geq W_1. \quad (10)$$

Since σ is Nash stable, no deviation is profitable. For every agent i , we know that $u_i(\sigma) \geq u_i(\sigma_{-i}, 3 - \sigma_i)$.⁶ Summing up this inequality over the set $\tilde{\mathcal{N}}$ of agents having a skill whose corresponding task is executed once under σ , we get an inequality $W_A \geq W_B$. W_A is the sum of the utilities of the agents in $\tilde{\mathcal{N}}$ under σ , and W_B is the sum of the utilities of the agents in $\tilde{\mathcal{N}}$ if they unilaterally deviate from σ . Note that W_A is upper bounded by the total weight of the tasks executed under σ , i.e., $W_1 + W_2 \geq W_A$. Since we have excluded skills owned by a single agent, we know that the weight of the tasks executed once under σ is shared by at least two agents. It follows that $W_B \geq 2W_1$ because for each task executed once, at least two agents get its full weight if she unilaterally deviates. We obtain $W_1 + W_2 \geq W_A \geq W_B \geq 2W_1 \Leftrightarrow (10)$.

We have $SW(\sigma^*) \leq 2(W_1 + W_2)$ because a task is executed at most twice ($q = 2$), and $SW(\sigma) = W_1 + 2W_2$. Use (10) multiplied by $\frac{1}{2}$ to get that $SW(\sigma) = W_1 + \frac{3}{2}W_2 + \frac{1}{2}W_2 \geq \frac{3}{2}(W_1 + W_2)$. Therefore, the PoA $\frac{SW(\sigma^*)}{SW(\sigma)}$ is upper bounded by $\frac{2(W_1 + W_2)}{\frac{3}{2}(W_1 + W_2)} = \frac{4}{3}$.

To conclude, instances where the PoA is exactly $4/3$ exist. \square

7 Conclusion and Future Work

In this article, we have focused on hedonic skill games and analyzed the existence, efficiency, and computation of Nash stable outcomes. Deciding whether a general instance admits a Nash stable outcome is an NP-complete problem when $q \geq 3$. A Nash stable outcome exists for every singleton agent instance, and it can be computed in polynomial time, but a natural dynamics like BRD can cycle. Nash stable outcomes also exist in singleton task instances, and the game admits a potential, but it is unlikely that a polynomial algorithm can compute a Nash stable outcome.

Several research directions that are worth investigating arise from this work. For instance, what is the difficulty of deciding whether an instance of the hedonic skill game admits a Nash stable outcome when $q = 2$? What is the exact PoA for singleton agents instances when $q = n$ and $\tau = 2$? In many cases we focused on $q \in \{2, n\}$, but examining the same problems for any value of q leaves many open questions. A lot of problems related to the better response dynamics remain unresolved when a Nash stable outcome exists. We have seen that BRD always converges if the tasks are singleton, but it can cycle for singleton agents. Special singleton agents instances, especially when the number of coalitions is small (e.g., $q = 2$), deserve attention. Another interesting question deals with the convergence of the *best* response dynamics, in which every deviation is a best response instead of a better response.

We believe that it would also be useful to consider the ϵ -Nash stable outcomes, where $\epsilon \geq 1$. They are outcomes where no agent can improve her utility by a multiplicative factor strictly greater than ϵ . This is particularly interesting for the setting with non-singleton agents for which the existence of Nash stable outcomes is not guaranteed or, it is hard to compute (as in the case of singleton tasks).

We have adopted the classic utilitarian social welfare, however, other kinds of social welfare could be considered, like for instance, the egalitarian social welfare which is defined as the minimum utility among all the agents' utilities, and the Nash social welfare which is defined as the product of the agents' utilities. It would also be interesting to analyze other notions than Nash stability, e.g., core stability, strong Nash stability, individual stability, etc.

Finally, we think that it would be important to consider different rewarding schemes (i.e., alternative ways to define the utility of an agent) that can guarantee the existence of stable outcomes and possibly show a lower price of anarchy.

⁶Here we assume that $\sigma_i \in \{1, 2\}$, corresponding to the choice of coalition C_1 or coalition C_2 , and $3 - \sigma_i$ simply designates the other strategy (i.e., 1 when $\sigma_i = 2$, and 2 when $\sigma_i = 1$). Moreover, (σ_{-i}, x) is the standard notation for the strategy profile σ' where $\sigma'_j = \sigma_i$ when $j \neq i$, and $\sigma'_j = x$ when $j = i$.

Acknowledgment: Laurent Gourvès is supported by Agence Nationale de la Recherche (ANR), project THEMIS ANR-20-CE23-0018.

References

- [AdK11] Haris Aziz and Bart de Keijzer. Complexity of coalition structure generation. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 191–198, South Carolina, 2011. IFAAMAS.
- [AS16] Haris Aziz and Rahul Savani. Hedonic games. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, pages 356–376. Cambridge University Press, USA, 2016.
- [BFF⁺19] Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. Optimality and nash stability in additive separable generalized group activity selection problems. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 102–108, Macao, China, 2019. ijcai.org.
- [BMJK10] Yoram Bachrach, Reshef Meir, Kyomin Jung, and Pushmeet Kohli. Coalitional structure generation in skill games. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, pages 703–708, Washington D.C (USA), 2010. AAAI Press.
- [BPR13] Yoram Bachrach, David C. Parkes, and Jeffrey S. Rosenschein. Computing cooperative solution concepts in coalitional skill games. *Artificial Intelligence*, 204:1–21, 2013.
- [BR08] Yoram Bachrach and Jeffrey S. Rosenschein. Coalitional skill games. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1023–1030, South Carolina, 2008. IFAAMAS.
- [ÇKO21] Bugra Çaskurlu, Fatih Erdem Kizilkaya, and Berkehan Ozen. Hedonic expertise games. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT)*, volume 12885 of *Lecture Notes in Computer Science*, pages 314–328, Denmark, 2021. Springer.
- [DEK⁺18] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, and Gerhard J. Woeginger. Group activity selection problem with approval preferences. *Int. J. Game Theory*, 47(3):767–796, 2018.
- [DG80] J. H. Drèze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica*, 48(4):987–1003, 1980.
- [Haj06] Jana Hajdjková. Coalition formation games: a survey. *International Game Theory Review*, 8(4):613–641, 2006.
- [IPE17] Ayumi Igarashi, Dominik Peters, and Edith Elkind. Group activity selection on social networks. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, pages 565–571, California, USA, 2017. AAAI Press.
- [Kar72] Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103, Boston, MA, 1972. Plenum Press.
- [KP99] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, New York (USA), 1999. Springer.
- [LLT09] Theodoros Lappas, Kun Liu, and Evimaria Terzi. Finding a team of experts in social networks. In *Proceedings of the 15th ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, pages 467–476, New York City, 2009. ACM.

- [MS96] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [NRTV07] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [OSC⁺15] Tenda Okimoto, Nicolas Schwind, Maxime Clement, Tony Ribeiro, Katsumi Inoue, and Pierre Marquis. How to form a task-oriented robust team. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 395–403, Richland, SC, 2015. ACM.
- [Ros73] Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [SHB⁺11] Walid Saad, Zhu Han, Tamer Basar, Mérouane Debbah, and Are Hjørungnes. Hedonic coalition formation for distributed task allocation among wireless agents. *IEEE Transactions on Mobile Computing*, 10(9):1327–1344, 2011.
- [SY91] Alejandro A. Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1):56–87, 1991.
- [vS21] Bernhard von Stengel. *Game Theory Basics*. Cambridge University Press, Cambridge (UK), 2021.
- [WD06] Michael J. Wooldridge and Paul E. Dunne. On the computational complexity of coalitional resource games. *Artificial Intelligence*, 170(10):835–871, 2006.