

Stability and Efficiency in Hedonic Project Games

Jaber Valizadeh , Dongmo Zhang and Omar Mubin

School of Computing, Data and Mathematical Sciences, Western Sydney University, Australia

{J.Valizadeh, D.Zhang, O.Mubin}@westernsydney.edu.au

Abstract

We introduce *Hedonic Project Games*, a model in which agents choose projects with divisible rewards while holding subjective preferences over coalition composition. This framework addresses the absence from existing models: agents care simultaneously about *who* they collaborate with and *what* they work on. We study three stability notions: classical *Nash Stability* and two refinements, *Joining Stability* and *Leaving Stability*, which account for the welfare of both the deviating agent and the affected coalition members. We evaluate the efficiency of stable outcomes using the *Price of Anarchy* and *Price of Stability*, comparing the social welfare of stable outcomes to that of an optimal allocation. While stable outcomes may not exist in general, we identify broad and natural preference classes in which stability and efficiency improve significantly. In particular, under monotonic decreasing preferences with coalition size, Nash and joining stability coincide and are guaranteed to exist, while leaving stability may fail. Under per-capita non-decreasing preferences, socially optimal outcomes are always Nash stable and coincide with leaving stability, although equilibrium inefficiency remains unbounded. Experiments on synthetic and real-world data support the theoretical efficiency results.

1 Introduction

It is the beginning of a new semester in a university research department. The head of the department introduces a set of m new research projects, each associated with its own funding and rewards, and invites n researchers, each of whom can join exactly one project. As you enter the meeting room, you see that several researchers have already chosen some projects. You would like to choose a project that maximizes your own share of the reward while also considering the colleagues involved: you may prefer to collaborate with close friends or perhaps reconnect with colleagues you have not worked with recently. This scenario captures a fundamental tension absent from most existing models: agents care simultaneously about *what* they work on and *with* whom they work. Your

decision depends not just on the project’s reward, but on how that reward is shared and whether your teammates enhance or diminish your experience. What should you choose? And when everyone reasons similarly, what outcomes emerge, and are they efficient?

In many multi-agent environments, agents face a dual and fundamentally conflicting incentive structure. On the one hand, they compete over objective, divisible resources (funding, rewards, or values), whose value is diluted by congestion and rival participation [Valizadeh and Telikani, 2026; Valizadeh *et al.*, 2025a]. On the other hand, they cooperate within coalitions whose composition directly affects individual satisfaction, productivity, or compatibility through subjective preferences over teammates. Existing models typically capture only one side of this tension. The *congestion games* [Rosenthal, 1973], *project games* [Bilò *et al.*, 2023], *market sharing games* [Goemans *et al.*, 2004] model competition over rewards, abstracting away from who one collaborates with and implicitly treating teammates as interchangeable. *Hedonic games* [Dreze and Greenberg, 1980; Bogomolnaia and Jackson, 2002], in contrast, model preferences over coalition composition but ignore the presence of external, rivalrous rewards tied to specific tasks or projects. *Group activity selection problems* [Darmann *et al.*, 2012] enable agents to choose both groups and activities; however, activities are typically labeled without tangible payoffs. As a result, neither framework adequately represents environments in which agents must simultaneously decide what to work on and with whom, balancing the benefits of collaboration against the costs of competition for limited resources.

We introduce *Hedonic Project Games* (HPGs), a framework that integrates subjective coalition preferences with objective, divisible project rewards. Each agent selects exactly one project; the resulting coalition determines both the agent’s share of the project reward and the subjective value they derive from their teammates. Formally, an agent’s utility is the product of their preference for the coalition and their equal share of the project’s reward. This multiplicative structure creates trade-offs that neither hedonic games nor project games capture alone: joining a high-reward project may be unappealing if it means undesirable colleagues or diluted rewards; a lower-reward project may be preferable with a compatible team; an empty project guarantees full reward but sacrifices collaboration benefits.

A central solution concept in both hedonic and project-based coalition formation is *Nash stability*, which requires that no agent can benefit by unilaterally deviating. However, Nash stability evaluates deviations solely from the deviating agent’s perspective, ignoring the incentives of agents directly affected by entry or exit. In our setting, where coalition composition and reward division are tightly coupled, this limitation leads to outcomes that are individually stable yet fragile at the coalition level. To address this, we introduce two refinements of Nash stability: *Joining Stability*, which accounts for the consent of agents in the target coalition, and *Leaving Stability*, which considers the impact of departure on the remaining coalition.

Stable outcomes may not exist under arbitrary preferences, as conflicting incentives may induce improvement cycles. We therefore identify two natural and expressive preference classes under which stability is guaranteed: *monotonically decreasing preferences* with coalition size, and *per-capita preferences*, where utility depends on average coalition value. For these classes, we establish existence results for Nash, joining, and leaving stable outcomes. We further analyze the efficiency of stable outcomes using the *Price of Anarchy* and *Price of Stability*, showing how preference structure fundamentally shapes equilibrium efficiency. Finally, we validate our theoretical findings through experiments on synthetic instances and real-world crowdsourcing data, demonstrating that the interaction between coalition preferences and reward dilution is both theoretically and empirically consequential.

2 Related Work

Hedonic games have been extensively studied [Banerjee *et al.*, 2001; Cechlárová and Romero-Medina, 2001; Bogomolnaia and Jackson, 2002], with a comprehensive survey provided by Aziz *et al.* [Aziz and Savani, 2016]. Several subclasses capture different preference structures, including Additively Separable Hedonic Games (ASHGs) [Bogomolnaia and Jackson, 2002], Fractional Hedonic Games (FHGs) [Aziz *et al.*, 2019], and Hedonic Coalition Nets [Elkind and Wooldridge, 2009]. Work on effective representations [Cechlárová and Romero-Medina, 2001; Ballester, 2004; Elkind and Wooldridge, 2009; Aziz *et al.*, 2019], stable outcome existence [Aziz and Brandl, 2012; Bilò *et al.*, 2018; Bullinger and Romen, 2024], and algorithmic aspects [Brandt *et al.*, 2021] is extensive. In ASHG, pure Nash equilibria exist under symmetric preferences but can be highly inefficient, with computation being PLS-complete. FHGs [Bilò *et al.*, 2018] guarantee Nash stability under restricted weights but have an unbounded price of anarchy. However, classical hedonic games ignore external factors such as task rewards, and even expressive variants assume agents only care about group composition.

Group Activity Selection Problems (GASPs) [Darmann *et al.*, 2012] generalize hedonic games by letting agents choose both a group and an activity. For o-GASP with strict preferences, deciding the existence of a Nash stable assignment is NP-complete [Darmann, 2015], with parameterized complexity studied in [Lee and Williams, 2017]. However, activities are usually labels without explicit pay-

offs. Project Games [Bilò *et al.*, 2023], a class of *monotone valid utility games* [Vetta, 2002; Augustine *et al.*, 2015; Goemans *et al.*, 2004; Marden and Roughgarden, 2010; Marden and Wierman, 2013], model agents choosing projects with fixed weight contributions. They connect to load balancing [Vöcking, 2007], congestion games [Rosenthal, 1973], and hedonic games [Dreze and Greenberg, 1980; Aziz and Savani, 2016]. These games admit pure Nash equilibria and have price of anarchy at most 2 [Goemans *et al.*, 2004], but assume agents are indifferent to teammates.

The most closely related model is AS-GGASP [Bilò *et al.*, 2019], but it differs fundamentally in utility structure. AS-GGASP assumes additive separability between coalition preferences and activity values, whereas hedonic project games use multiplicative utilities coupling coalition composition with divisible rewards. This captures interactions additive models cannot express: preferred collaborators matter only when sufficient reward remains, and high-reward projects become less attractive when shared among unsuitable coalitions. Unlike AS-GGASP’s fixed activity utilities, hedonic project games explicitly model competition over divisible rewards. Each agent’s utility is its normalized coalition preference multiplied by an equal share of the project reward, creating an endogenous trade-off between coalition attractiveness and reward dilution. Normalization preserves comparability without changing the multiplicative, size-dependent structure. Consequently, utilities in hedonic project games cannot be decomposed into additive agent-agent and agent-project components.

3 Definitions and Notation

We consider a finite set of agents $N = \{1, 2, \dots, n\}$, and a finite set of projects $M = \{j_1, j_2, \dots, j_m\}$, where each agent $i \in N$ can perform any project $j \in M$. Each project $j \in M$ has an associated positive reward defined by the reward function $r : M \rightarrow \mathbb{R}_+$, where $r(j)$ represents the reward for project j . Each agent $i \in N$ must select exactly one of the m projects. We denote the project chosen by agent i as their *strategy*, $s_i \in M$. A *strategy profile* (also referred to as an *outcome* throughout the remainder of this paper) is a vector $\mathbf{s} = (s_1, \dots, s_n) \in \mathcal{S}$, where the strategy space is $\mathcal{S} = M^n$.

Let $\mathcal{N}_i = \{C \subseteq N : i \in C\}$ denote the set of possible coalitions that each player $i \in N$ belongs to. A hedonic project game is formally defined as follows:

Definition 1. A hedonic project game is specified by $G = \langle N, M, r, \mathcal{S}, (p_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where:

- $N = \{1, 2, \dots, n\}$ is the set of n players.
- $M = \{j_1, j_2, \dots, j_m\}$ is a finite set of m projects.
- $r : M \rightarrow \mathbb{R}_+$ is a reward function.
- $\mathcal{S} = M^n$ denotes the strategy space, where M is the strategy space for each player i .
- $p_i : \mathcal{N}_i \rightarrow [0, 1]$ denotes player i ’s preference function¹, where $\sum_{C \in \mathcal{N}_i} p_i(C) = 1$.
- $u_i : \mathcal{S} \mapsto \mathbb{R}_+$ denotes the utility function for each player $i \in N$.

Given a strategy profile $\mathbf{s} \in \mathcal{S}$, the *coalition* for project $j \in M$ is defined as $C_j(\mathbf{s}) = \{i \in N : s_i = j\}$, which is the set of players that select project j under $\mathbf{s}$, and we write $|C_j(\mathbf{s})|$ for its size. The reward $r(j)$ of each project $j \in M$ is divisible and is shared equally among all players in $C_j(\mathbf{s})$.

Given a strategy profile $\mathbf{s} \in \mathcal{S}$, the utility of each player $i \in N$ is defined as follows:

$$u_i(\mathbf{s}) = p_i(C_{s_i}(\mathbf{s})) \cdot \frac{r(s_i)}{|C_{s_i}(\mathbf{s})|} \quad (1)$$

where $p_i(C_{s_i}(\mathbf{s}))$ is player i 's preference for the coalition to which she belongs, $r(s_i)$ is the reward of the chosen project s_i , and $|C_{s_i}(\mathbf{s})|$ is the number of players selecting project s_i .

3.1 Connections to Existing Models

We briefly show that hedonic project games generalize several well-studied models under natural restrictions on rewards and preferences.

- Identical Project Rewards: All projects yield the same reward, normalized to 1, i.e., $\forall j \in M, r(j) = r$. Rewards are *generic* (heterogeneous) when this does not hold.

- Identical Preferences: Each agent assigns equal value to all coalitions they belong to (indifferent among coalitions), i.e., $\forall i \in N, \forall C, C' \in \mathcal{N}_i, p_i(C) = p_i(C')$.

- Additively Separable Preferences: Let $w_i : N \rightarrow [0, 1]$ be the *valuation function* of an agent i assigning a weight $w_i(\ell) \in [0, 1]$ being in the same coalition as i . The preferences of players in hedonic project games are *additivity separable*, i.e., the preference of each agent i for every coalition $C \in \mathcal{N}_i$ is $p_i(C) = \sum_{\ell \in C} w_i(\ell)$, assume $w_i(i) = 0$.

Hedonic project games with identical preferences reduced to *singleton congestion games* [Rosenthal, 1973] and their variants such as *symmetric project games with identical agents* [Bilò et al., 2023], *market-sharing games* [Goemans et al., 2004]. Load-balancing games with identical machines [Vöcking, 2007] correspond to hedonic project games with identical preferences and rewards. Hedonic project games with additively separable preferences and identical project rewards reduce to fractional hedonic games with non-negative preferences [Aziz et al., 2019; Bilò et al., 2019].

3.2 Stability Notions

The classical stability notion in hedonic games is *Nash stability* (NS), named after the Nash equilibrium [Aziz and Savani, 2016]. It guarantees that no agent has an incentive to unilaterally change their coalition, regardless of the consent of others.

In our setting, a *deviation* occurs when a player i changes her strategy from the current project s_i under the strategy profile $\mathbf{s}$ to a distinct project $j \neq s_i$. In hedonic project games, a strategy profile $\mathbf{s}$ is *Nash stable* if, for all players $i \in N$ and all $j \in M$, it holds

$$u_i(\mathbf{s}) \geq u_i(j, \mathbf{s}_{-i})$$

where $(j, \mathbf{s}_{-i})$ denotes the profile obtained from $\mathbf{s}$, after player i deviates from s_i to j . The set of all Nash stable outcomes for the game G is denoted by $\text{NS}(G)$.

The limitation of Nash stability is that it considers only the deviator's utility, which can be overly restrictive in cooperative settings. To address this, we introduce two refinements of Nash stability that explicitly model coalition-level consent constraints.

Definition 2 (Joining Stability (JS)). *A strategy profile $\mathbf{s}$ is joining stable if it is Nash stable, and for all $i \in N, j \in M$, and $\ell \in C_j(\mathbf{s})$,*

$$u_\ell(\mathbf{s}) \geq u_\ell(j, \mathbf{s}_{-i})$$

where $C_j(\mathbf{s})$ denotes the set of players who selected the project j under strategy profile $\mathbf{s}$. The set of all joining stable outcomes for the game G is denoted by $\text{JS}(G)$.

Definition 3 (Leaving Stability (LS)). *A strategy profile $\mathbf{s}$ is leaving stable if for all $i \in N, \ell \in C_{s_i}(\mathbf{s})$, and $j \in M$,*

$$u_\ell(\mathbf{s}) \geq u_\ell(j, \mathbf{s}_{-i})$$

where $C_{s_i}(\mathbf{s})$ denotes the set of players who selected the same project as player i under strategy profile $\mathbf{s}$. The set of all leaving stable outcomes for the game G is denoted by $\text{LS}(G)$.

Our notions of joining and leaving stability strengthen Nash stability by imposing stricter conditions on how deviations impact the involved coalitions. In contrast, Individual Stability (IS) and Contractual Nash Stability (CNS) only partially account for affected members' preferences, allowing deviations that can still produce fragile outcomes. IS blocks a deviation only if someone in the target coalition is worse off, which may trap agents in suboptimal coalitions. CNS adds consent requirements but still does not guarantee stability for all affected members.

However, joining stability ensures that no unilateral deviation benefits the deviator, and additionally that no member of the target coalition would benefit from the deviator joining. Similarly, leaving stability ensures that no unilateral deviation benefits the deviator and that no member of the current coalition would benefit from the deviator leaving. Overall, joining stability and leaving stability are stronger than Nash stability, while IS and CNS are weaker. These stricter conditions ensure that coalitions have no incentive to actively reshape their membership, making joining stability and leaving stability stronger refinements of Nash stability than IS or CNS.

The logical hierarchy among these stability concepts is illustrated in Figure 1. Arrows indicate implication relationships; for example, either joining stability or leaving stability implies Nash stability.




Figure 1: Logical relationships between our solution concepts.

¹Normalization ensures comparability across coalitions while preserving the full expressiveness of agents' preference orderings.

3.3 Social Welfare and Efficiency Measures

Utilitarian social welfare (commonly referred to simply as social welfare) represents the collective well-being of all participants (players) in a game, typically measured by aggregating individual utilities or payoffs. Given a strategy profile $\mathbf{s} \in \mathcal{S}$, the social welfare is defined by

$$\text{SW}(\mathbf{s}) = \sum_{i \in N} u_i(\mathbf{s}) \quad (2)$$

A *social optimum* of a game G is a strategy profile $\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{S}} \text{SW}(\mathbf{s})$ that maximizes the social welfare. We denote by $\text{OPT}(G) = \text{SW}(\mathbf{s}^*)$ the corresponding value.

Given a game instance, the Price of Anarchy (PoA) [Koutsoupias and Papadimitriou, 1999] is defined as the ratio of the worst case objective function value of a stable outcome to that of an optimal one. Analogously, the Price of Stability (PoS) [Anshelevich *et al.*, 2008] captures the best-case ratio between the social optimum and the welfare at a stable outcome.

Definition 4. Fix a game G , and $\alpha \in \{\text{NS}, \text{JS}, \text{LS}\}$. Let $\alpha(G)$ denote the set of α -stable outcomes for G . For any outcome $\mathbf{s}$, let $\text{SW}(\mathbf{s})$ be its social welfare and $\text{OPT}(G)$ the maximum welfare. The price of anarchy and the price of stability under α , are

$$\text{PoA}_\alpha(G) = \max_{\mathbf{s} \in \alpha(G)} \frac{\text{OPT}(G)}{\text{SW}(\mathbf{s})}, \quad (3)$$

$$\text{PoS}_\alpha(G) = \min_{\mathbf{s} \in \alpha(G)} \frac{\text{OPT}(G)}{\text{SW}(\mathbf{s})} \quad (4)$$

In both cases, if $\text{SW}(\mathbf{s}) = 0$ for some $\mathbf{s} \in \alpha(G)$, the corresponding ratio is defined as ∞ . We consider the PoA and PoS to be *undefined* if no stable outcome exists with respect to any of the stability notions.

Let $\mathcal{G}$ be a class of hedonic project games, and $\alpha \in \{\text{NS}, \text{JS}, \text{LS}\}$. Define $\text{PoA}_\alpha(\mathcal{G}) = \sup_{G \in \mathcal{G}} \text{PoA}_\alpha(G)$ and $\text{PoS}_\alpha(\mathcal{G}) = \sup_{G \in \mathcal{G}} \text{PoS}_\alpha(G)$. By definition, the following bounds hold:

$$\begin{aligned} \text{PoA}_{\text{JS}}(\mathcal{G}), \text{PoA}_{\text{LS}}(\mathcal{G}) &\leq \text{PoA}_{\text{NS}}(\mathcal{G}), \\ \text{PoS}_{\text{JS}}(\mathcal{G}), \text{PoS}_{\text{LS}}(\mathcal{G}) &\geq \text{PoS}_{\text{NS}}(\mathcal{G}) \end{aligned}$$

To this end, we now focus on several natural subclasses of preferences that have been widely studied in the literature on hedonic games, and examine their implications within the context of the hedonic project game setting.

3.4 Structured Preferences

Instead of allowing arbitrary preferences that give rise to instability and inefficiency, we focus on structured preferences where systematic rules guide agents' choices: (i) monotonic decreasing preferences for coalition size, and (ii) per-capita preferences.

- (i). **Monotonic Decreasing Preferences:** The preferences of players in hedonic project games are *monotonic decreasing* with respect to coalition size, when for all $i \in N$, and $C, C' \in \mathcal{N}_i$, holds:

$$p_i(C) \geq p_i(C') \iff |C| \leq |C'|$$

Preferences are *agent-specific* monotonic decreasing if for each agent $i \in N$, there exists a function $f_i : \mathbb{N} \rightarrow [0, 1]$ with $f_i(k) > f_i(k+1)$ for all $k \geq 1$, such that for all coalitions $C \in \mathcal{N}_i$, $p_i(C) = f_i(|C|)$; if $f_i = f$ for all i , they are *universal*.

- (ii). **Per-Capita Preferences:** In a hedonic project game, for each player $i \in N$ and each coalition $C \in \mathcal{N}_i$, the *per-capita preference* of player i in coalition C is defined as

$$\hat{p}_i(C) = \frac{p_i(C)}{|C|}$$

Player i is said to have *per-capita (monotonic) non-decreasing* (respectively, *non-increasing*) preferences if for any two coalitions $C, C' \in \mathcal{N}_i$ with $C' \subset C$, it holds that

$$\hat{p}_i(C) \geq \hat{p}_i(C') \quad (\text{respectively, } \hat{p}_i(C) \leq \hat{p}_i(C'))$$

Preferences are called *per-capita strictly decreasing* (respectively, *strictly increasing*) if the above inequality is strict for all such pairs $C' \subset C$. Per-capita monotonicity only constrains preferences over nested coalitions; no assumptions are made for incomparable coalitions.

Monotonic decreasing preferences (in coalition size) capture a strict form of size aversion: agents always prefer smaller coalitions, and any increase in size weakly decreases utility, independent of who joins or how productive the coalition becomes. In contrast, per-capita preferences evaluate a coalition based on the average benefit per member. Under this view, adding a new agent may increase or decrease utility depending on whether that agent raises or lowers the coalition's average payoff.

Intuitively, monotonic decreasing preferences reflect congestion or crowding effects driven purely by group size, while per-capita preferences capture efficiency considerations where growth is beneficial only if the new member improves average performance. For example, in interdisciplinary research teams, a biologist joining computer scientists may significantly improve the project's effectiveness, increasing the average benefit rather than merely adding another member.

4 Existence and Efficiency of Stable Outcomes

In this section, we analyze the existence, and efficiency of stable outcomes in hedonic project games based on Nash stability and its coalition-consent refinements, joining and leaving stability under structured preference.

4.1 General Existence and Efficiency

In hedonic project games, a Nash stable outcome does not necessarily exist under *unrestricted* preferences (see Example 1. in the appendix²). Moreover, even when a Nash stable outcome is guaranteed to exist, the corresponding PoA can still be *unbounded* (see Example 2 in the appendix).

We show that the inefficiency of Nash stable outcomes (if they exist) remains controlled in hedonic project games. For any game G , let $r_{\max} = \max_{j \in M} r(j)$ and $r_{\min} =$

²The appendix is provided in this link.

$\min_{j \in M} r(j)$, denote the largest and smallest project rewards, respectively. We say that for a game G , rewards are ρ -bounded if the ratio $\frac{r_{\max}}{r_{\min}} \leq \rho$ expresses the degree of reward disparity, where ρ is a constant. Preferences are δ -bounded if $p_i(C) \geq \delta$ for all $i \in N$, and $C \in \mathcal{N}_i$. The following theorem establishes a constant-factor upper bound on the PoA in terms of ρ and δ .

Theorem 1. *Let $G \in \mathcal{G}$ be a hedonic project game that admits at least one Nash stable outcome. If rewards are ρ -bounded, and preferences are δ -bounded, then we get*

$$\text{PoA}_{\text{NS}}(G) \leq \frac{\rho}{\delta}$$

Proof. Let $\mathbf{s}$ be an arbitrary Nash-stable outcome of G , and let $\mathbf{s}^*$ be a strategy profile maximizing social welfare, with value $\text{OPT}(G)$.

We first derive an upper bound on $\text{OPT}(G)$. For any outcome $\mathbf{s}$ and any project $j \in M$ with coalition $C_j(\mathbf{s})$, the contribution of j to social welfare is

$$\sum_{i \in C_j(\mathbf{s})} p_i(C_j(\mathbf{s})) \cdot \frac{r(j)}{|C_j(\mathbf{s})|} = \frac{r(j)}{|C_j(\mathbf{s})|} \sum_{i \in C_j(\mathbf{s})} p_i(C_j(\mathbf{s}))$$

Since preferences are normalized, $p_i(C) \leq 1$ for all i and all $C \in \mathcal{N}_i$, implying

$$\sum_{i \in C_j(\mathbf{s})} p_i(C_j(\mathbf{s})) \leq |C_j(\mathbf{s})|$$

Hence, the welfare contribution of project j is at most $r(j)$.

By ρ -bounded we have $r_{\max} \leq \rho \cdot r_{\min}$. As at most $\min(n, m)$ projects can be non-empty, we obtain

$$\text{OPT}(G) \leq \min(n, m) \cdot r_{\max} \leq \min(n, m) \cdot \rho \cdot r_{\min} \quad (5)$$

We now lower-bound the social welfare $\text{SW}(\mathbf{s})$.

Case 1: There exists an empty project j . For any agent $i \in N$, Nash stability implies

$$u_i(\mathbf{s}) \geq u_i(j, \mathbf{s}_{-i}) = p_i(\{i\}) \cdot r(j) \geq \delta \cdot r_{\min}$$

Summing over all agents yields

$$\text{SW}(\mathbf{s}) \geq n \cdot \delta \cdot r_{\min} \quad (6)$$

Case 2: No project is empty (i.e., all m projects are used). Note that this is only possible if $m \leq n$, as there are only n agents. For each project $j \in M$,

$$\begin{aligned} \sum_{i \in C_j(\mathbf{s})} p_i(C_j(\mathbf{s})) \cdot \frac{r(j)}{|C_j(\mathbf{s})|} \\ = r(j) \left(\frac{1}{|C_j(\mathbf{s})|} \sum_{i \in C_j(\mathbf{s})} p_i(C_j(\mathbf{s})) \right) \end{aligned} \quad (7)$$

By δ -boundedness, $p_i(C_j(\mathbf{s})) \geq \delta$ for all $i \in C_j(\mathbf{s})$, implying that the term in parentheses is at least δ . Thus,

$$\text{SW}(\mathbf{s}) \geq \delta \sum_{j \in M} r(j) \geq \delta \cdot m \cdot r_{\min} \quad (8)$$

In Case 1, by combining (5) and (6), we obtain

$$\text{PoA}_{\text{NS}}(G) \leq \frac{n \cdot \rho \cdot r_{\min}}{n \cdot \delta \cdot r_{\min}} = \frac{\rho}{\delta}$$

In Case 2, since $m \leq n$, by combining (5) and (8), we obtain

$$\text{PoA}_{\text{NS}}(G) \leq \frac{m \cdot \rho \cdot r_{\min}}{\delta \cdot m \cdot r_{\min}} = \frac{\rho}{\delta}$$

Since $\mathbf{s}$ was arbitrary, the bound holds for all Nash-stable outcomes. Moreover, δ -bounded preferences and positive rewards ensure $\text{SW}(\mathbf{s}) > 0$. $\square$

4.2 Existence and Efficiency under Structured Preferences

Non-existence of stable outcomes arises from preference-induced deviation cycles that prevent convergence to equilibrium. We therefore focus on structured preference classes that guarantee equilibrium existence.

Additively Separable Preferences

A hedonic project game with additively separable preferences and identical project rewards reduces to a fractional hedonic game with non-negative weights. Existing results show that in additively separable hedonic games with non-negative preferences, the grand coalition is trivially Nash stable, since no agent can improve by deviating to a singleton coalition (e.g., [Olsen, 2009; Aziz and Savani, 2016; Bilò *et al.*, 2018; Valizadeh *et al.*, 2025b]). Observation 1 extends this result to hedonic project games: even with project-specific rewards, the grand coalition remains stable and is both joining and leaving stable.

Observation 1. *Let $G \in \mathcal{G}$ be a hedonic project game with additively separable preferences. If all players select the same project under a strategy profile $\mathbf{s}$, then $\mathbf{s}$ is a Nash, joining and leaving stable outcome.*

While the existence of a Nash, joining and leaving stable outcome is guaranteed for hedonic project games with additively separable preferences, the PoA for this subclass of preferences can be arbitrarily large.

Theorem 2. *Let $\mathcal{G}$ be a class of hedonic project games with additively separable preferences. Then $\text{PoA}_{\text{NS}}(\mathcal{G})$ is unbounded.*

Monotonic Decreasing Preferences

Under *universally* monotonic decreasing preferences, every hedonic project game admits an exact potential function and therefore converges to a Nash-stable outcome.³ In contrast, under *agent-specific* monotonic decreasing preferences, a potential function may not exist, but a Nash-stable outcome is still guaranteed. These games are player-specific congestion games with decreasing payoffs and satisfy the Finite Improvement Property (FIP), implying existence via Milchtaich [Milchtaich, 1996].

For general monotonic decreasing preferences in coalition size, Nash and joining stability coincide, while leaving stability may fail to exist (see Example 3 in the appendix). When project rewards are identical, equilibrium coalitions are size-balanced, and every Nash-stable outcome is socially optimal.

³Such games are a subclass of congestion games [Rosenthal, 1973].

Observation 2. Given a game instance $G \in \mathcal{G}$, if preferences are monotonic decreasing (w.r.t. coalition size), then $NS(G) = JS(G)$.

In hedonic project games, when the project rewards are identical and the agents' preferences decrease monotonically with coalition size, Nash stability ensures a strong balance across coalition sizes. The following lemma formalizes this property, demonstrating that at equilibrium, no coalition can be significantly larger than others.

Lemma 1. Let $G \in \mathcal{G}$ be a hedonic project game with monotonic decreasing preferences. If project rewards are identical and $\mathbf{s}$ is a Nash stable or joining stable outcome of G , then the sizes of the coalitions induced by $\mathbf{s}$ differ by at most one.

We say that monotonic decreasing preferences are *concave*, if the function $f_i : \mathbb{N} \rightarrow [0, 1]$, satisfies

$$f_i(k+2) - 2f_i(k+1) + f_i(k) \leq 0 \quad \text{for all } k \geq 1$$

where $p_i(C) = f_i(|C|)$ for every agent $i \in N$ and coalition $C \in \mathcal{N}_i$.

Theorem 3. Let $G \in \mathcal{G}$ be a hedonic project game with universal concave monotonic decreasing preferences. If project rewards are identical, then $PoS_{NS}(G) = PoS_{JS}(G) = 1$.

Per-Capita Preferences

Under per-capita non-decreasing preferences, trivially socially optimal allocations are Nash stable outcomes.

Observation 3. Let $G \in \mathcal{G}$ be a hedonic project game with per-capita non-decreasing preferences. Let $j^* \in \arg \max_{j \in M} r(j)$, and let $\mathbf{s}^*$ be the strategy profile in which all agents select project j^* . Then $\mathbf{s}^*$ is a Nash stable outcome and an optimal allocation.

While leaving and joining stability are generally refinements of Nash stability, the following result shows that under per-capita non-decreasing preferences, these two notions coincide.

Theorem 4. Let $G \in \mathcal{G}$ be a game instance with per-capita non-decreasing preferences. Then the sets of Nash and leaving stable outcomes coincide, i.e.,

$$NS(G) = LS(G)$$

Proof. We show that $NS(G) \subseteq LS(G)$, as the reverse inclusion $LS(G) \subseteq NS(G)$ holds by definition.

Let $\mathbf{s}$ be an arbitrary Nash stable outcome. Fix an agent $i \in N$ and consider a deviation from project s_i to some project $j \in M$.

If the deviation is strictly unprofitable for agent i , i.e., $u_i(\mathbf{s}) > u_i(j, \mathbf{s}_{-i})$, then the leaving stability condition is trivially satisfied. Hence, it suffices to consider the case where agent i is indifferent:

$$u_i(\mathbf{s}) = u_i(j, \mathbf{s}_{-i})$$

Case 1: If $|C_{s_i}(\mathbf{s})| = 1$, then agent i forms a singleton coalition at project s_i , and $C_{s_i}(\mathbf{s}) \setminus \{i\} = \emptyset$. In this case, there are no other agents in i 's original coalition whose utility could be affected by i 's departure. The leaving stability condition requires that for all $\ell \in C_{s_i}(\mathbf{s})$, we have $u_\ell(\mathbf{s}) \geq u_\ell(j, \mathbf{s}_{-i})$.

Since $C_{s_i}(\mathbf{s}) = \{i\}$ and agent i is indifferent (or the deviation is unprofitable), the LS condition is trivially satisfied for all $\ell \neq i$ and directly for $\ell = i$.

Case 2: If $|C_{s_i}(\mathbf{s})| \geq 2$, let $\ell \in C_{s_i}(\mathbf{s}) \setminus \{i\}$ be any agent remaining in project s_i after i leaves. Their utilities before and after the deviation are

$$u_\ell(\mathbf{s}) = \frac{p_\ell(C_{s_i}(\mathbf{s})) \cdot r(s_i)}{|C_{s_i}(\mathbf{s})|}$$

and

$$u_\ell(j, \mathbf{s}_{-i}) = \frac{p_\ell(C_{s_i}(\mathbf{s}) \setminus \{i\}) \cdot r(s_i)}{|C_{s_i}(\mathbf{s})| - 1}$$

respectively.

The leaving stability requires $u_\ell(\mathbf{s}) \geq u_\ell(j, \mathbf{s}_{-i})$, which simplifies to:

$$\frac{p_\ell(C_{s_i}(\mathbf{s}))}{|C_{s_i}(\mathbf{s})|} \geq \frac{p_\ell(C_{s_i}(\mathbf{s}) \setminus \{i\})}{|C_{s_i}(\mathbf{s})| - 1}$$

after canceling $r(s_i) > 0$. This is precisely

$$\hat{p}_\ell(C_{s_i}(\mathbf{s})) \geq \hat{p}_\ell(C_{s_i}(\mathbf{s}) \setminus \{i\})$$

This inequality holds by the assumption that preferences are per-capita non-decreasing.

Since the argument applies to every agent i and every indifferent deviation, $\mathbf{s}$ satisfies the leaving stability condition and hence $NS(G) \subseteq LS(G)$. Thus,

$$NS(G) = LS(G)$$

□

Corollary 1. Let $G \in \mathcal{G}$ be a hedonic project game. If preferences are per-capita non-decreasing, then

$$PoS_{NS}(G) = PoS_{LS}(G) = 1$$

While leaving stability is guaranteed to exist in the hedonic project game with non-decreasing preferences and coincide with Nash stability, Nash stability does not imply joining stability (see Example 4 in the appendix).

Under per-capita non-increasing preferences, each project can be viewed as a resource whose payoff to an agent depends only on the number of agents selecting it and is weakly decreasing in congestion. Consequently, hedonic project games with per-capita non-increasing preferences form a class of player-specific congestion games with decreasing payoff functions, which are known to satisfy the finite improvement property [Milchtaich, 1996].

Theorem 5. Let $G \in \mathcal{G}$ be a game instance with per-capita non-increasing preferences. Then

$$NS(G) = JS(G)$$

The following result establishes that per-capita strictly decreasing preferences preclude leaving stability for any non-trivial coalition structure.

Observation 4. For hedonic project games with per-capita strictly decreasing preferences, no outcome containing a coalition of size at least two can be leaving stable.

Corollary 2. Under per-capita strictly decreasing preferences, the only possible leaving stable outcomes are those where every coalition is a singleton. Consequently, if $n > m$, no leaving stable outcome exists.

5 Experimental Evaluation

We conduct an experimental evaluation to validate the theoretical results using both real-world and synthetic datasets⁴. For a real-world dataset, we have used a publicly available dataset from *gMission*, a research-oriented and open-source crowdsourcing platform [Chen *et al.*, 2014], which comprises 1200 tasks collected from real-world crowdsourcing activities, each characterized by publish time, geographic location, and associated rewards [Tong *et al.*, 2019]. We analyze 1,500 instances drawn from 10 *gMission* datasets. We infer agents’ preferences from *gMission* by defining them as a decreasing function of distance between agents, with closer agents interpreted as having stronger collaborative affinity.

Synthetic instances follow a multi-scale design: small-scale ($n \sim \mathcal{U}(2, 7)$, $m \sim \mathcal{U}(2, 10)$) to validate exact methods; medium-scale ($n \sim \mathcal{U}(8, 15)$, $m \sim \mathcal{U}(11, 20)$) for scalability and approximation analysis; and large-scale ($n \sim \mathcal{U}(16, 25)$, $m \sim \mathcal{U}(21, 30)$) to assess computational limits. Project rewards follow three models: *uniform* $\mathcal{U}(5, 100)$; *Gaussian* $\mathcal{N}(52.5, \lambda^2)$ with λ chosen so that 99.7% of values lie in $[5, 100]$ and clipped accordingly; and *exponential* with mean 50, shifted and clipped to $[5, 100]$. For each reward model, we generate 100 independent instances, yielding 300 synthetic games.

For each player i , preferences are normalized so that $\sum_{C \ni i} p_i(C) = 1$. Identical preferences assign equal value to all coalitions. Additively separable preferences use a symmetric weight matrix $w_i(j)$ with $w_i(j) \sim U[0.1, 1]$, defining $p_i(C) = \sum_{j \in C \setminus \{i\}} w_i(j)$. Monotonic decreasing and per-capita preferences are size-based, $p_i(C) = f_i(|C|)$, derived from a base function $g(k)$ enforcing the required monotonicity (harmonic for monotonic decreasing, quadratic for per-capita non-decreasing, and constant for per-capita non-increasing), and normalized using the exact number of coalitions of each size. All randomness uses a fixed seed.

5.1 Results and Analysis

Figure 2 shows clear separation in PoA across preference classes. Identical preferences exhibit low inefficiency (mean 1.09, max 1.45), with mass concentrated near 1, consistent with the bound $\text{PoS}_{\text{NS}}(\mathcal{G}) \leq 2$ for singleton congestion games. Per-capita non-decreasing preferences perform best (mean 1.03, median 1.00), with most instances achieving optimal welfare. Monotonic decreasing preferences yield the highest inefficiency (mean 1.55, max 3.32) and largest variance, reflecting fragmentation induced by aversion to large coalitions. Per-capita non-increasing preferences show intermediate inefficiency (mean 1.42, max 2.06) with a roughly symmetric distribution.

Figure 3 reveals that reward heterogeneity significantly impacts equilibrium efficiency: the exponential distribution exhibits the highest PoA variance due to *star* projects that attract crowding, while uniform and Gaussian distributions yield tighter, more moderate PoA/PoS values as their homogeneous rewards reduce coordination sensitivity.

⁴The code is provided in this link.




Figure 2: PoA behavior across preference classes.




Figure 3: Comparing the PoA and PoS across the three reward distributions.

6 Conclusion

We introduced and analyzed *Hedonic Project Games*, a strategic framework that captures the joint dependence of agents’ outcomes on both coalition composition and project rewards. The results show that when coalition preferences and reward sharing are tightly coupled, stability and efficiency are driven mainly by how agents evaluate coalitions rather than by reward magnitude. Stability increases when agents account for how their participation affects others, even indirectly. Preference structures that scale with coalition size reduce incentives for disruptive deviations, while unrestricted preferences enable gains from reward dilution without internalizing its broader effects. This suggests that inefficiency in collaborative settings is caused more by misaligned preference scaling than by competition over rewards alone.

This work also points to several future directions, including heterogeneous reward-sharing schemes, relaxing the single-project assumption, and studying learning dynamics in repeated hedonic project games with limited feedback. From an algorithmic perspective, the complexity of finding stable outcomes under more general preference classes remains open. More broadly, hedonic project games provide a flexible model for collaborative systems where agents jointly choose both tasks and partners, supporting the design and analysis of modern multi-agent platforms.

Acknowledgements

This work was supported by the Commonwealth through an Australian Government Research Training Program Scholarship [DOI: <https://doi.org/10.82133/C42F-K220>].

References

- [Anshelevich *et al.*, 2008] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [Augustine *et al.*, 2015] John Augustine, Ning Chen, Edith Elkind, Angelo Fanelli, Nick Gravin, and Dmitry Shiryayev. Dynamics of profit-sharing games. *Internet Mathematics*, 11(1):1–22, 2015.
- [Aziz and Brandl, 2012] Haris Aziz and Florian Brandl. Existence of stability in hedonic coalition formation games. *arXiv preprint arXiv:1201.4754*, 2012.
- [Aziz and Savani, 2016] Haris Aziz and Rahul Savani. Hedonic games. 2016.
- [Aziz *et al.*, 2019] Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. Fractional hedonic games. *ACM Transactions on Economics and Computation (TEAC)*, 7(2):1–29, 2019.
- [Ballester, 2004] Coralio Ballester. Np-completeness in hedonic games. *Games and Economic Behavior*, 49(1):1–30, 2004.
- [Banerjee *et al.*, 2001] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18(1):135–153, 2001.
- [Bilò *et al.*, 2018] Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. Nash stable outcomes in fractional hedonic games: Existence, efficiency and computation. *Journal of Artificial Intelligence Research*, 62:315–371, 2018.
- [Bilò *et al.*, 2019] Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. Optimality and nash stability in additively separable generalized group activity selection problems. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 102–108, 2019.
- [Bilò *et al.*, 2023] Vittorio Bilò, Laurent Gourvès, and Jérôme Monnot. Project games. *Theoretical Computer Science*, 940:97–111, 2023.
- [Bogomolnaia and Jackson, 2002] A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- [Brandt *et al.*, 2021] Felix Brandt, Martin Bullinger, and Anaëlle Wilczynski. Reaching individually stable coalition structures in hedonic games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5211–5218, 2021.
- [Bullinger and Romen, 2024] Martin Bullinger and René Romen. Stability in online coalition formation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9537–9545, 2024.
- [Cechlárová and Romero-Medina, 2001] Katarína Cechlárová and Antonio Romero-Medina. Stability in coalition formation games. *International Journal of Game Theory*, 29:487–494, 2001.
- [Chen *et al.*, 2014] Zhao Chen, Rui Fu, Ziyuan Zhao, Zheng Liu, Leihao Xia, Lei Chen, Peng Cheng, Caleb Chen Cao, Yongxin Tong, and Chen Jason Zhang. gmission: A general spatial crowdsourcing platform. *Proceedings of the VLDB Endowment*, 7(13):1629–1632, 2014.
- [Darmann *et al.*, 2012] Andreas Darmann, Edith Elkind, Sascha Kurz, Jérôme Lang, Joachim Schauer, and Gerhard Woeginger. Group activity selection problem. In *Internet and Network Economics: 8th International Workshop, WINE 2012, Liverpool, UK, December 10-12, 2012. Proceedings 8*, pages 156–169. Springer, 2012.
- [Darmann, 2015] Andreas Darmann. Group activity selection from ordinal preferences. In *International Conference on Algorithmic Decision Theory*, pages 35–51. Springer, 2015.
- [Dreze and Greenberg, 1980] Jacques H Dreze and Joseph Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society*, pages 987–1003, 1980.
- [Elkind and Wooldridge, 2009] Edith Elkind and Michael J Wooldridge. Hedonic coalition nets. In *AAMAS (1)*, pages 417–424. Citeseer, 2009.
- [Goemans *et al.*, 2004] Michel Goemans, Li Erran Li, Vahab S Mirrokni, and Marina Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing*, pages 55–66, 2004.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, pages 404–413. Springer, 1999.
- [Lee and Williams, 2017] Hooyeon Lee and Virginia Vassilevska Williams. Parameterized complexity of group activity selection. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 353–361, 2017.
- [Marden and Roughgarden, 2010] Jason R Marden and Tim Roughgarden. Generalized efficiency bounds in distributed resource allocation. In *49th IEEE Conference on Decision and Control (CDC)*, pages 2233–2238. IEEE, 2010.
- [Marden and Wierman, 2013] Jason R Marden and Adam Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.

- [Milchtaich, 1996] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.
- [Olsen, 2009] Martin Olsen. Nash stability in additively separable hedonic games and community structures. *Theory of Computing Systems*, 45(4):917–925, 2009.
- [Rosenthal, 1973] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [Tong *et al.*, 2019] Yongxin Tong, Yuxiang Zeng, Bolin Ding, Libin Wang, and Lei Chen. Two-sided online micro-task assignment in spatial crowdsourcing. *IEEE Transactions on Knowledge and Data Engineering*, 33(5):2295–2309, 2019.
- [Valizadeh and Telikani, 2026] Jaber Valizadeh and Ray Telikani. Online learning for decentralized multi-agent planning in repeated hedonic skill games. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 36, pages 342–350, 2026.
- [Valizadeh *et al.*, 2025a] Jaber Valizadeh, Dongmo Zhang, and Omar Mubin. Autonomy with structural task allocation games: From inefficiency to optimality. In *International Conference on Principles and Practice of Multi-Agent Systems*, pages 435–452. Springer, 2025.
- [Valizadeh *et al.*, 2025b] Jaber Valizadeh, Dongmo Zhang, and Omar Mubin. Learning preferences in additive separable hedonic project games. In *Australasian Joint Conference on Artificial Intelligence*, pages 491–505. Springer, 2025.
- [Vetta, 2002] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 416–425. IEEE, 2002.
- [Vöcking, 2007] Berthold Vöcking. Selfish load balancing. *Algorithmic game theory*, 20:517–542, 2007.