Minimax Preferences for Tiered Coalition Formation Games

Nicholas Fluty[a,*], **Judy Goldsmith**[a] **and Brent Harrison**[a]

[a]University of Kentucky Dept of Computer Science
ORCID (Nicholas Fluty): https://orcid.org/0009-0007-3250-4915, ORCID (Judy Goldsmith): https://orcid.org/0000-0002-8383-5390, ORCID (Brent Harrison): https://orcid.org/0000-0002-1301-5928

Abstract.

Tiered coalition formation games (TCFGs) are coalition formation games in which the coalitions are ordered, changing the nature of agent preferences. TCFGs have been primarily used to model problems in gaming and competitive scenarios, where a tier list may effectively represent the various strengths of agents in intransitive relationships. Past work on TCFGs introduced preference frameworks that determine preference profiles from matchup data, but they rely on assumptions that may not align with real-world applications.

To address these limitations, we introduce *minimax preferences*, a new framework that considers von Neumann's notion of optimal minimax strategies. We first define our framework and prove its distinctness from previous TCFG frameworks. When we apply previous notions of stability, we find that our framework guarantees the existence of a unique Nash stable tier list. We prove that this tier list is also uniquely core stable, making it a strong general solution to the tier list formation problem. We provide a polynomial-time algorithm that generates this tier list given a matrix that defines agent matchups. Lastly, we provide an example of a stable tier list and analyze the output of our algorithm on randomized matrices.

1 Introduction

Coalition formation games (CFGs) [12, 9, 6] explore the problem of allocating agents into groups (coalitions) based on their preferences between different allocations (coalition structures). We can say that an agent's preference is determined by what other agents it "sees" in each, which is strictly the set of agents in its coalition in the case of a *hedonic* CFG [8]. *Tiered coalition formation games* (TCFGs) [15] vary from hedonic games by introducing an ordering to the coalitions from bottom to top tier, and an agent sees the set of agents in its tier and below. The outcome of a TCFG is a *tier list*, which can be considered stable (through various notions of stability) when no agent or group of agents can see a more preferred set of agents by making an allowed change to the allocation.

Tier lists can be commonly seen in the context of competitive intransitive games, such as in the Pokémon and Super Smash Bros. series of games. Rather than simply ranking characters from best to worst, tier lists place character options of similar strength or importance in the same tier, which may better represent intransitive relationships. Tier lists tend to primarily consider usage and win rates, but there seems to be no standard methodology for constructing one from given matchup data.

To better model these matchup-based games, Siler introduced the *matchup-oriented preferences* framework for TCFGs [15]. This allows agent preferences to be represented through a relation or matrix, where agent pairs (matchups) either have a winner or result in a tie. Under this framework, an agent strictly prefers to see sets of agents against which it has more wins and fewer losses. This will generally cause an agent to prefer a tier list where it has more good matchups against agents in lower tiers and bad matchups against agents in higher tiers.

While Siler's representation handles deterministic matchups between agents, Arnold, *et al.*, state that we may often want to rank the strength of agents that instead have probabilistic matchup outcomes [2]. To account for this, they extend the framework, calling it *probabilistic* matchup-oriented preferences, where preferences are defined using a win probability matrix. This matrix contains expected matchup outcomes, or payoffs in $[-1, 1]$, so an agent's preference is for sets where its sum of payoffs is greater.

In order for a tier list assignment to be considered *Nash stable*, there must be no agent that can move itself to a different existing tier or newly inserted tier where it sees a more preferred set of agents [15]. Consider an example where an agent loses more than it wins against other agents in its current tier; that agent would gain utility by moving down a tier, so the tier list is not Nash stable. Due to the zero-sum nature of matchup outcomes, Nash stability requires that all agents within the same tier on average tie against each other (sum of zero payoff), else one would want to move down a tier.

When we consider randomized probabilistic matrices, sums of exactly zero are almost impossible, so Nash stability is unlikely to occur with multiple agents in the same tier. In fact, Siler proved that every TCFG with matchup-oriented preferences has at least one Nash stable tier list consisting strictly of singleton tiers [15], which is effectively just a total ranking of individual agents. This is unfortunate because Nash stability is sometimes taken to indicate the desirability of a coalition assignment.

This concern has inspired additional work to make it possible to have stable tier lists with a higher percentage of tiers being non-singleton, referred to as a tier list's *friendliness* [2]. Arnold, *et al.*, introduced the notions of *socially conscious stability* [4] and ϵ-*stability* [1], which we discuss in the next section. By using these new notions of stability, friendlier stable tier lists became possible under matchup-oriented preferences. We note that, while doing this

* Corresponding Author. Email: nickfluty39@gmail.com

gives the tier list construction problem more solutions, it also has the apparent side effect of making agents settle for less preferable tier assignments.

The main weakness we find with matchup-oriented preferences is that the matchup outcomes are summed with equal weights, as if opposing players select them with equal probability. This assumption does not necessarily hold for our problem, as the existence of a tier list suggests that some characters may be better and thus more popular than others. We should instead expect that, when all players select optimally, there is an underlying probability distribution over agents in use. Thus, our problem is to find a probability distribution (or distributions) that can achieve equilibrium when used by all players, and then make agent preferences properly weight their matchups according to those probabilities. Because matchup outcomes are entirely based on a zero-sum matrix, this reduces the problem to a matrix game, which opens a better means for exploring optimal strategies.

Von Neumann discusses that two players in a matrix game can find optimality by using *minimax strategies* [17]. A minimax strategy is defined as a probability vector for selecting an action (row or column) that minimizes the maximum payoff (chances of winning) of any opposing strategy. When applied to a symmetric game (meaning both players have the same actions), the minimum expected payoff of a minimax strategy is zero, meaning that it is unbeatable and ties against itself. Thus, when both players use the minimax strategy, neither player has any incentive to use a different strategy, creating Nash equilibrium. Calculating this strategy effectively solves the game, removing any need to guess the opponent's strategy.

Applying minimax strategies to a group of agents allows us to identify which ones are theoretically most important to the game. Specifically, we consider the *support* of a minimax strategy, defined as the set of agents chosen with non-zero probability. We use this to formally define *minimax preferences* for TCFGs, which, given an input matrix of probabilistic matchup outcomes, produce preference profiles for agents based on minimax strategies that can be calculated for each group. Because it would be disadvantageous for players to select agents not in the support of a minimax strategy, we say that agents strictly prefer groups where they are in the support of a minimax strategy. This is defined more formally in Section 4.

After introducing our framework, we show that a given payoff matrix always produces different preference profiles under probabilistic matchup-oriented preferences and minimax preferences. Furthermore, there are preference profiles that can be represented by matchup-oriented preferences and not minimax, and vice-versa. We then investigate the relationships of different stability notions when TCFGs are defined by minimax preferences.

We show that preference profiles produced through our framework have the special property of having exactly one tier list that is Nash stable for those agents. Because this tier list is unique, we call it the *minimax tier list*. We briefly discuss how the friendliness of the minimax tier list may change based on the relative strengths of agents in the payoff matrix, thus creating a starting point for its potential analytical usages. To make the minimax tier list more accessible, we clearly define an algorithm for finding it, and we prove that our algorithm runs in polynomial time.

The remainder of the paper is organized as follows. In Section 2, we provide definitions related to TCFGs, matchup-oriented preferences, and various notions of stability. In Section 3, we review literature related to this work. In Section 4, we detail minimax preferences, and we prove that the minimax preference framework encodes different preferences than other TCFG frameworks do. We also investigate stability notions for minimax TCFGs and analyze the complexity of finding a stable tier list. In Section 5, we give an example matrix and resulting minimax tier list, and we briefly assess the averaged results of a larger set of randomized matrices. Finally, Section 6 concludes the paper.

2 Definitions

In this section, we provide formal definitions of TCFGs, previous preference frameworks, and notions of stability that we discuss in this paper. We introduce new definitions associated with our contribution later in Section 4.

2.1 Tiered Coalition Formation Games

We first provide a definition of tiered coalition formation games:

Definition 1. *[15] A* tiered coalition formation game (TCFG) *instance is a pair $(A, \succeq)$ consisting of a finite set of agents $A = \{a_1, a_2, \ldots, a_n\}$ and a preference profile $\succeq = \{\succeq_1, \succeq_2, \ldots, \succeq_n\}$, where $\succeq_i$ is a weak total order over subsets of A that contain a_i.*

An outcome (tier list) consists of a partition of A into a totally ordered set of k disjoint coalitions (tiers) $T = \{T_1, T_2, \ldots, T_k\}$, where we say T_1 is the lowest tier.

In a tier list T, we say that a_i sees *a_j if a_i is in same tier as a_j or a higher tier: Seen$(a_i, T) = \bigcup_{l=1}^{m} T_l$ for $a_i \in T_m$. For tier lists T and T', $\succeq_i$ specifies a_i's preference between Seen(a_i, T) and Seen(a_i, T').*

Each agent needs to have associated with it a set of preferences that can be used to identify, for any two sets of agents, which it prefers to see. Note that Seen(a_i, T) is the set of agents seen by agent a_i from its position in the tier list T.

While tiers are commonly represented in order from lowest to highest as in $\{T_1, ..., T_k\}$, we often use in this paper an equivalent tier list in reverse order as in $\{T_{\text{top}}, ..., T_{\text{top}-k+1}\}$. We use T_{top} rather than T_k because this allows us to refer to the top tiers of incomplete tier lists, as well as the top tiers of multiple tier lists with different values for k. In all cases, we let this translation from $\{T_{\text{top}}, ..., T_{\text{top}-k+1}\}$ to $\{T_1, ..., T_k\}$ be implicit.

2.2 Matchup-Oriented Preferences

Siler introduced *matchup-oriented preferences* to TCFGs so that agent preferences could be represented through a simple matchup model [15]. This is done through an antisymmetric, antireflexive relation, where each pair defines a matchup as being "good" for one agent and "bad" for the other. If there is not a pair containing two given agents, the matchup between them is considered a tie. An agent's preference in this framework is for a tier list that maximizes its *utility*, calculated as the count of good matchups it sees (agents see those in the same tier or below) minus the count of bad matchups it sees.

Siler's matchup-oriented preferences treat matchups as deterministic, not allowing for varying strengths of favorability. Arnold, *et al.*, [2] extended this definition to allow for win *probabilities,* discussing that this model is more versatile while maintaining the fundamental properties of its predecessor. We provide a formal definition of the extended framework in terms of a payoff matrix:

Definition 2. *[2]* Probabilistic matchup-oriented preferences *for a TCFG are defined by an $n \times n$ zero-sum payoff matrix M. An agent*

a_i has a preference to maximize its utility*, calculated as the sum of its payoffs against the agents it sees, i.e.,*

$$S \succ_i S'$$

iff

$$\sum_{a_j \in S} M[i,j] > \sum_{a_j \in S'} M[i,j].$$

For both of these preference frameworks, an agent's utility is the sum of its preferences toward the agents it sees. It is important to note that the utility is an equally-weighted sum, so all outcomes are viewed as equally important to an agent. This makes sense when each agent represents a single entity, meaning that each truly sees one of everything else, or if agents have matchups against all other agents in equal proportion as in a round-robin tournament. As previously discussed, this is often not the case for intransitive games, which motivates us to create a framework that may better model these problems.

2.3 Stability Notions

We use *stability* to assess the accuracy or desirability of a tier list based on the potential for improving changes. When no improving changes can be made, the tier list is stable. Different notions of stability assess different types of changes and different conditions for considering them to be improving. We are especially concerned with Nash stability:

Definition 3. *[15] A tier list T is* Nash stable *if there exists no move (to an existing tier, or by forming a new tier in the hierarchy) by any individual agent a_i that would yield a new tier list T' such that*

$$\mathit{Seen}(a_i, T') \succ_i \mathit{Seen}(a_i, T).$$

Nash stability is desired in many social situations because it means that each agent sees its situation as ideal when considering unilateral moves.

In order for a tier list to be considered Nash stable under either of the matchup-oriented preference frameworks, there must be no agent that can increase its utility by unilaterally moving to a different or new tier. When we consider the summing nature of utility, we only need to consider the net gain of utility from the agents it begins to see after moving up. Likewise, we only need to consider the net loss of utility from the agents it no longer sees after moving down. For example, if the sum of an agent's payoffs within its current tier is negative, we know it should desire to move down to exclude them from its set of seen agents and thus increase its utility.

We also consider *core stability*:

Definition 4. *[15] A tier list T is* core stable *if there exists no nonempty set of agents $B \subset A$ that could form a new tier together such that, for the resulting tier list T', it holds that*

$$\mathit{Seen}(a_i, T') \succsim_i \mathit{Seen}(a_i, T) \quad \text{for all } a_i \in B,$$

with the preference being strict for at least one $a_i \in B$.

Core stability addresses the possibility of a set of agents moving to form a new tier. If at least one agent in the set strictly prefers this movement, while the others in the set consider it to be at least as good as the current assignment, the tier list is not core stable.

Other notions of stability for TCFGs include *socially conscious stability* [4] and *ϵ-stability* [1]. We do not provide formal definitions of these because they do not directly apply to our minimax preferences framework, as they are defined based on a utility value. Effectively, socially conscious stability is the equivalent of contractually individual stability in hedonic games [5].

Arnold, *et al.*'s ϵ-stability introduces a threshold to another notion of stability (e.g. ϵ-Nash stability), where changes that result in a utility gains of less than ϵ do not block stability.

While these notions are designed to show tier list accuracy or desirability in the absence of Nash stability, we designed minimax preferences so a Nash stable tier list can always be found, and it only has a singleton tier when an agent beats all other agents below it.

3 Related Work

In this section, we discuss related works in the areas of hedonic games and TCFGs. We also detail what we consider most relevant from game theory concerning minimax strategies and their importance.

3.1 Gaming Related Hedonic Games

This work falls into the category of coalition formation games [12, 9, 6], specifically coalition formation games without transferrable utility. The original work was inspired by gaming considerations, as was work on other games [16, 15, 14, 2, 11, 3] that are not necessarily tiered. While some of that related work models preferences as hedonic [8], i.e., agents having preferences only over their own coalitions, the work to TCFGs is not hedonic. Agents in TCFGs base their preferences not only on their own tier, but on those below. If we considered all seen agents, we could also view TCFGs as an overlapping coalition formation game [7], though that is not an approach we take here. In general, work on coalition formation games often focuses, as we do here, on various notions of stability or optimality. Work includes proofs of the existence (or possible nonexistence) of stable coalition structures, the computational complexity of deciding whether stable coalition structures exist or of finding stable or optimal coalition structures, as well as coalition formation algorithms and heuristics.

3.2 Tiered Coalition Formation Games

In 2017, Siler introduced tiered coalition formation games (TCFGs) along with the deterministic matchup-oriented preference framework [15]. Representing preferences in this way was shown to give TCFGs various desirable properties, such as the equivalence of Nash stability with core stability and the guaranteed existence of a Nash stable tier list. While a Nash stable solution to the TCFG problem could be found in polynomial time, the only solution guaranteed to exist consisted of strictly singleton sets.

Regarding general TCFGs, Siler proved that the problem of deciding whether there exists a Nash stable tier list for a given TCFG is NP-complete, and deciding whether there exists a core stable tier list is NP-hard [15]. It should be noted that matchup-oriented preferences and our minimax preferences both guarantee that Nash and core stable tier list exist.

Generalizing for any TCFG preferences, Siler proved that core stability implies Nash stability. Under matchup-oriented preferences, we are given that the reverse is also true, that Nash stability implies core stability [15]. Since we later show that our minimax preference framework has this same property, we generally discuss Nash stability with an implication of this equivalence.

Because matchup outcomes are often probabilistic, Arnold, *et al.*, formalized probabilistic matchup-oriented preferences, which define preferences based on a matrix of expected matchup outcomes between agents [2]. Using this extension was shown to maintain those previously mentioned properties, while it also provides the versatility for more complex intransitive games.

Arnold, *et al.*, [2] also introduced k-TCFGs, a variant of TCFGs that restricts tier lists to at most k tiers. This problem was inspired particularly by Smogon's tier list for the game of Pokémon. Because each of Smogon's tiers corresponds to a rule set banning the use of Pokémon above it, it makes sense to be able to limit the number of tiers for the purpose of their application. Because singleton tier lists are not possible when $k < n$, k-TCFGs under matchup-oriented preferences have been shown to not always have Nash stable solutions.

Because Nash stable solutions were often either nonexistent or consisting only of singleton tiers, more recent work has focused on defining new notions of stability that consider a smaller subset of possible moves. Arnold, *et al.*, introduced socially conscious stability so that k-TCFGs would always have at least one stable solution [4]. While this is interesting for k-TCFGs, this notion is equivalent to Nash stability and core stability for general TCFGs under matchup-oriented preferences.

Lastly, Arnold, *et al.*, consider a recent empirical evaluation of the renowned Smogon tier list for Pokémon battling [3]. Using observed win probabilities for battles between individual Pokémon, this work showed that the Smogon tier list is not stable by any notions described in this paper. Note, however, that the Smogon tier list is not designed for individual battles but teams, which affects the interpretation of those results.

In the case of Pokémon battles, we note that this is a game where players choose their agents based on their varying advantages in matchups, which is likely not in equal proportions. Because modeling this type of problem seems to be the focus of past work, we show how we can use minimax strategies to assess which agents are most optimal to select against other intelligent players.

The next section covers a background on minimax strategies and their relevance to this problem.

3.3 Minimax Strategies

Von Neumann introduced modern game theory, which explores optimal strategies in two-person zero-sum matrix games [17]. In this model, the first and second players simultaneously select actions i and j, respectively, to received the expected payoffs of $M[i, j]$ and $-M[i, j]$, respectively, as defined by a matrix M. The objective in a matrix game is to deploy a *strategy*, defined as a probability vector over actions, that maximizes one's expected payoff against their opponent's strategy.

Von Neumann discussed the importance of using a *minimax strategy*, which is a strategy that maximizes your expected payoff while assuming a worst-case scenario where your opponent's strategy minimizes your payoff. At the core of game theory, von Neumann's minimax theorem says that the maximum of a player's minimum expected payoffs equals the minimum of their opponent's maximum expected losses [17].

This theorem is critical to game theory because it proves that the performance of a minimax strategy can always be minimized by an opposing minimax strategy. Neither player is incentivized to do anything other than use a minimax strategy, creating Nash equilibrium. A pair of minimax strategies is considered a *solution* to the matrix game.

Calculating a minimax strategy can be done in polynomial time using a polynomial-time linear programming solver, such as Karmakar's algorithm [13]. Considering that most known practical applications of matrix games have relatively small action spaces, it is easy to find this solution.

Previous work in TCFGs has not yet discussed matrix games or any form of strategies over agents. We show in our work that, by modeling agents in a TCFG as the actions in a matrix game, we are able to better analyze the usefulness of each agent when selected optimally by players. Our case is a *symmetric game*, meaning both players have the same set of actions (agents) and payoffs, causing the minimax strategies to be the same with an average payoff of zero for both players.

The agents selected within a minimax strategy, which we say are in the *support* of the minimax strategy, have special importance to the solved game. Players have a mathematical reason to select each of them some portion of the time; that is, if a player is known by their opponent to not use a minimax strategy, there is a strategy that can be used to beat them over half of the time.

We also know that there is an average payoff of zero when the minimax strategy is used against a pure strategy of any agent in its support. If this were not so, the minimax strategy would either lose to itself or have an agent in its support that beats it, neither of which can be true. Thus, even though some agents should be selected less often than others, they can be regarded as equally strong when they are used. This serves as a justification for considering all of these agents to be the "best", which we use to later define their preferences toward groups.

4 Minimax Preferences for TCFGs

Recall that matchup-oriented preferences, in terms of the resulting preference profiles, make the most sense under the assumption that agents see all matchups in equal proportions. However, outside of round-robin tournaments, this assumption may not hold well, making the preferences less meaningful. Consider the case where agents represent selectable characters in a game, which is to say that agents can be likened to selectable actions in a modified rock-paper-scissors; rather than expecting a round-robin among characters, we can expect that players will select characters with different probabilities, typically based on their perceived strength.

As an extreme example, consider the case where agent a wins with 100% probability against all agents except agent b, to which it loses more often than it wins. Suppose that agent b loses to all other agents 100% of the time, with the exception of agent a. If two players should select an agent to determine which one of them wins (like one does in rock-paper-scissors), where should agent b be ranked in a tier list for this game? Even though agent b significantly loses to the majority of other agents, it should be highly ranked for its singular role in countering agent a. Despite this, matchup-oriented preferences would yield that agent b should prefer being at the bottom of the tier list, where its total utility is zero rather than negative.

To address these limitations, we introduce *minimax preferences* for TCFGs. Our approach is goal-oriented, where our goal is to consistently produce a meaningful Nash stable tier list. To achieve this, we design the preference profiles so that, for any group of agents, only the "best" agents in the group prefer to see the entire group. By making this preference strictly stronger toward larger groups, we assure that Nash stable tier lists will always maximize how high each agent is placed. To define what agents should be considered best, or

most preferable to a player, we apply the concept of optimal minimax strategies from game theory.

In the following subsections, we first define minimax preferences for TCFGs. We then show that minimax preference profiles are distinct from matchup-oriented preference profiles. We then investigate interesting properties of stability that can be observed under our framework. Lastly, we introduce and analyze an algorithm that can be used to find Nash stable tier lists.

4.1 Definition of Minimax Preferences

Definition 5. Minimax preferences *for a TCFG are defined by an $n \times n$ zero-sum payoff matrix M. We say that a* strategy *is an n-dimension probability vector that assigns weights to agents, and the* mixed strategy space Δ_S *contains all strategies that assign weights strictly to agents contained in the set S, i.e.,*

$$\Delta_S = \left\{ x \in [0,1]^n \;\middle|\; \sum_{i=1}^{n} x_i = 1 \text{ and } \forall\, a_i \notin S \;\; [x_i = 0] \right\}.$$

We say that a minimax strategy *of S is a strategy in Δ_S that minimizes the maximum expected loss (equally, maximizes the minimum expected payoff) against all strategies in Δ_S. The space containing all minimax strategies of S is denoted by Δ_S^*, i.e.,*

$$\Delta_S^* = \left\{ y^* \in \Delta_S \;\middle|\; y^* = \arg\min_{y \in \Delta_S} \max_{x \in \Delta_S} x^\top M y \right\}.$$

An agent's preference is to maximize the number of agents it sees while being in the support of a minimax strategy, i.e.,

$$S \succ_i S'$$

iff

$$\exists\, y^* \in \Delta_S^* \text{ such that } y_i^* > 0$$

and either

$$\forall\, y^* \in \Delta_{S'}^* \; [y_i^* = 0]$$

or

$$|S| > |S'|.$$

For any group of agents, there is a minimax strategy. We say that an agent is in the *support* of a minimax strategy if its probability therein is greater than zero. By definition, any agent in the support of a group prefers that group over any other group where it is not in the support. If it is in the support of both groups, it prefers the bigger one. For any set of agents, there is at least one agent that prefers to see the complete set. We also note that a singleton set has only one strategy, which means its one agent is always in the support of a minimax strategy.

Minimax preferences account for the possible existence of multiple minimax strategies for a set of agents, such as when there are multiple Condorcet winners (they beat all other agents but tie against each other). While this may not occur in many games, minimax preferences define that all agents in the support of at least one minimax strategy of a group prefer to see that group.

Because agents prefer to see bigger groups, Nash stability may be achieved only when agents are in the highest tier in which they are in the support of a minimax strategy. Similarly, agents currently in the support of a minimax strategy never prefer to move downward. Thus, if agents are greedily allocated from the top tier down, we can reason that a Nash stable solution is possible. In fact, there will always exist exactly one Nash stable tier list, which can be found in polynomial time. We formally prove this and provide our algorithm for finding it in Section 4.4.

4.2 Preference Profile Differences

Minimax preferences and probabilistic matchup-oriented preferences are both defined by a payoff matrix, but we show that the resulting agent preferences are always different. After this, we compare the space of preference profiles that can be produced by any matrix.

Theorem 1. *There are no payoff matrices that produce the same set of preference profiles under both probabilistic matchup-oriented preferences and minimax preferences for any $n \geq 3$.*

Proof. This is broken into two cases:

Case 1: There does not exist a Condorcet winner, that is, for every agent a_i there exists some agent a_j such that $M[i,j] < 0$. Under probabilistic matchup-oriented preferences, this means that a_i prefers the set $A \setminus \{j\}$ over the set A, so no agent prefers the set A over all other sets. Under minimax preferences, there always exists some agent a_i that is in the support of a minimax strategy of A. Because A is larger than all of its proper subsets, a_i prefers the set A over all other sets.

Case 2: There is a Condorcet winner, that is, there exists some agent a_i such that, for every agent a_j, $M[i,j] \geq 0$. This is broken into two sub-cases:

Case 2.1: There is at least one tie between agents, that is, $M[i,j] = 0$ for some $i \neq j$. Under probabilistic matchup-oriented preferences, a_i has no preference between the sets $\{a_i\}$ and $\{a_i, a_j\}$. Under minimax preferences, a_i prefers the set $\{a_i, a_j\}$ over the set $\{a_i\}$.

Case 2.2: There are no ties between agents, that is, $M[i,j] \neq 0$ for all $i \neq j$. There must exist two agents a_j and a_k where $i \neq j \neq k$ and $M[j,k] > 0$. Under probabilistic matchup-oriented preferences, a_j prefers the set $\{a_i, a_j, a_k\}$ over the set $\{a_i, a_j\}$ because including a_k increases its utility. Under minimax preferences, the Condorcet winner a_i is the only agent in the minimax strategy of either set. Thus a_j has no preference between the sets $\{a_i, a_j, a_k\}$ and $\{a_i, a_j\}$. □

While minimax preferences and probabilistic matchup-oriented preferences always yield different preference profiles for a given matrix, we note that it is still possible for a tier list to be Nash stable under both frameworks. A notable example of this is a singleton ranking for a fully transitive game, where we know that every subset of agents has a Condorcet winner. While matchup-oriented preferences specify that agents prefer groups where they lose less over where they lose more, minimax preferences specify that agents are indifferent when they are not Condorcet winners. Despite this, these differences do not concern the agents in the stable tier list assignment because none of them see a set where they lose.

Next, we examine the space of preference profiles possible in each framework. We prove that each framework can produce profiles that cannot be produced by the other, so neither framework is a special case of the other.

Theorem 2. *Not all preference profiles that are possible under probabilistic matchup-oriented preferences are possible under minimax preferences.*

Proof. Consider a preference profile produced by the zero matrix under probabilistic matchup-oriented preferences. Every agent is indifferent towards all sets because the sum of its payoffs is always zero.

This is not possible under minimax preferences because at least one agent is part of a minimax strategy of A as a whole, which makes it prefer seeing all agents over seeing any proper subset of agents. □

Theorem 3. *Not all preference profiles that are possible under minimax preferences are possible under probabilistic matchup-oriented preferences.*

Proof. Consider a preference profile produced by the zero matrix under minimax preferences. Because the expected payoff of every pair of strategies is zero, every agent is in the support of a minimax strategy of every set, so an agent's preference is strictly towards the larger set. This is not possible under probabilistic matchup-oriented preferences because two agents can't prefer to see each other, as this would mean that the matchup has a positive payoff for both agents. □

4.3 Stability under Minimax Preferences

In this section, we consider stability under minimax preferences. We prove here that a Nash stable tier list is both guaranteed to exist and unique. We then prove that it is also core stable.

4.3.1 Existence of a unique Nash stable tier list

Theorem 4. *There exists exactly one Nash stable tier list for any TCFG under minimax preferences.*

Proof. For any Nash stable tier list assignment T over a set of agents A with minimax preferences, we first argue that there is a unique top tier it must have. We use this to show the uniqueness of each tier by induction.

For any $a \in A$, $\text{Seen}(a, T) = A$ iff $a \in T_{\text{top}}$. Let S_A be the set of agents in the support of at least one minimax strategy of A; all agents in S_A prefer to be in T_{top} over any other tier by definition of minimax preferences. Consider that any agent that moves to a new bottom tier $T_{\text{top}-k}$ will only see itself and thus be in the support of a minimax strategy; all agents not in S_A prefer to be in $T_{\text{top}-k}$ over T_{top} by definition of minimax preferences. Therefore, if $T_{\text{top}} \neq S_A$, T is not Nash stable.

Continuing recursively, consider the remaining set of agents $B = A \backslash S_A$ and the set of agents S_B in the support of at least one minimax strategy of B; by the same logic, if $T_{\text{top}-1} \neq S_B$, T is not Nash stable. This can be continued until the remaining set of agents is the empty set. Let $T^* = \{S_A, S_B, \ldots\}$; T^* is Nash stable because every agent is in the highest tier where it can be in the support of a minimax strategy. Since T is Nash stable iff $T = T^*$, there exists exactly one Nash stable tier list assignment under minimax preferences. □

Because minimax preferences give us a unique Nash stable tier list, we call it the *minimax tier list*. The most meaningful tier of the minimax tier list is T_{top} because it contains strictly the set of agents that players can select when playing optimally. In theory, all agents outside of the top tier are completely irrelevant to optimal play, so a game with only a few agents at its top tier may not see many different agents selected in practice.

Because optimality has already been addressed, lower tiers only receive their intended meaning when we apply the assumption that higher-tier agents may not be selected. Thus, the second tier reflects the optimality of a sub-game of the main game. When we consider the main game and all $k - 1$ sub-games, every agent has a game where it can be used optimally. In the case of a completely transitive game without ties, the minimax tier list will have n singleton tiers. Similarly, we expect that the average number of agents in each tier is generally greater for games with more intransitivity.

4.3.2 Equivalence of Nash and core stability

We first show that the minimax tier list is Nash stable, which means that Nash stability implies core stability under minimax preferences. Because the reverse was proven by Siler for all TCFGs, this makes them equivalent under our framework.

Theorem 5. *The minimax tier list is core stable*

Proof. Suppose a minimax tier list T is not core stable. Let B be a blocking tier for T. There exists a tier list T' that has B as a tier, where no agent in B strictly prefers T, and at least one agent in B strictly prefers T'.

We have previously shown that every agent in a minimax tier list is in the support of a minimax strategy of the set it currently sees, so there cannot be an agent in B that sees fewer agents in T' than in T. Thus, this movement cannot be downward for any agent.

Consider the set of agents seen from the new tier in T'. Because no agents have moved downward, this same set can be seen by any agent moving individually to that tier. Because at least one agent prefers to move here to see this set, T is not Nash stable. This contradicts that the minimax tier list is Nash stable.

Assuming that the minimax tier list is not core stable leads to a contradiction. □

Because the unique Nash stable tier list is also core stable, Nash stability implies core stability under minimax preferences. This observation can be combined with Siler's theorem so that we get the following:

Corollary 6. *Nash stability and core stability are equivalent under minimax preferences*

4.4 Finding the Minimax Tier List

As previously discussed, there is only one Nash stable tier list for a TCFG under minimax preferences, which we call the minimax tier list. Here we present an algorithm that shows how we can find the minimax tier list from an input matrix in polynomial time. While we account for the possibility of multiple minimax strategies by testing the viability of each unassigned agent with its own linear program, we note that there is often only one minimax strategy in practice, meaning that one linear program per tier may often suffice.

Observation 7. *Algorithm 1 runs in polynomial time.*

To analyze the time complexity of Algorithm 1, we first note that finding solutions to matrix games has been shown to be equivalent to linear programming by reduction [10]. Linear programming methods such as Karmakar's algorithm [13] run in polynomial time. At each new tier, we solve no more than n LPs. Because each tier will have at least one agent, this gives us a total number of LPs bounded by n^2, also a polynomial. Thus, the algorithm as a whole has polynomial time complexity.

5 Example Minimax Tier List

We provide an example minimax tier list outputted by Algorithm 1 using a randomized payoff matrix. After, we briefly discuss some averages resulting from running our algorithm on 100 other randomized matrices with the same specifications. This section is designed to show the feasibility of applying our minimax preference framework to find meaningful tier lists.

Algorithm 1 Minimax Tier List Construction

Input: Set of agents $A = \{a_1, a_2, \ldots, a_n\}$; skew-symmetric matrix $M \in [-1, 1]^{n \times n}$; minimum strategy weight $\epsilon > 0$
Output: Tier list $T = [T_{\text{top}}, T_{\text{top}-1}, \ldots, T_{\text{top}-k+1}]$ partitioning A

$U \leftarrow A$ {Unassigned agents}
$T \leftarrow [\,]$ {Initialize empty tier list}
$i \leftarrow 0$ {Tier index}
while $U \neq \emptyset$ **do**
  $T_{\text{top}-i} \leftarrow \emptyset$
  Let M_U be the submatrix of M restricted to agents in U
  for each $a_j \in U$ **do**
    Let $index(j)$ be the index of a_j in U
    Solve the following LP:

$$\begin{aligned} \text{maximize:} \quad & v \\ \text{subject to:} \quad & x^\top \cdot M_U \geq v \cdot \mathbf{1} \\ & x \geq \mathbf{0} \\ & \mathbf{1}^\top \cdot x = 1 \\ & x_{index(j)} \geq \epsilon \end{aligned}$$

    if $|v| \leq 10^{-6}$ **then**
      $T_{\text{top}-i} \leftarrow T_{\text{top}-i} \cup \{a_j\}$
    end if
  end for
  Append $T_{\text{top}-i}$ to T
  $U \leftarrow U \setminus T_{\text{top}-i}$
  $i \leftarrow i + 1$
end while
return T

For simplicity, we consider a case with 12 agents. We randomized the payoff matrix using a uniform distribution in the closed interval [-1.0,1.0], rounding to the nearest hundredth before running our algorithm. This matrix is shown in Table 1.

Table 1. Randomized Payoff Matrix

	a_1	a_2	a_3	a_4	a_5	a_6	a_7	a_8	a_9	a_{10}	a_{11}	a_{12}
a_1	0	.10	.43	.21	.09	-.15	.29	-.12	.78	.93	-.23	.58
a_2	-.10	0	.06	.14	.85	-.86	-.83	-.96	.67	.56	.74	.96
a_3	-.43	-.06	0	.60	-.08	.56	-.76	.28	-.71	.89	.04	-.17
a_4	-.21	-.14	-.60	0	-.47	.55	-.09	.14	-.96	.24	.22	.23
a_5	-.09	-.85	.08	.47	0	.89	.36	-.28	-.13	.40	-.88	.33
a_6	.15	.86	-.56	-.55	-.89	0	.34	-.58	-.74	-.37	-.27	.14
a_7	-.29	.83	.76	.09	-.36	-.34	0	-.12	.98	-.80	-.58	-.68
a_8	.12	.96	-.28	-.14	.28	.58	.12	0	.31	-.49	-.07	-.51
a_9	-.78	-.67	.71	.96	.13	.74	-.98	-.31	0	-.68	-.78	.31
a_{10}	-.93	-.56	-.89	-.24	-.40	.37	.80	.49	.68	0	-.72	-.61
a_{11}	.23	-.74	-.04	-.22	.88	.27	.58	.07	.78	.72	0	-.26
a_{12}	-.58	-.96	.17	-.23	-.33	-.14	.68	.51	-.31	.61	.26	0

Following the steps of Algorithm 1, we solved the full matrix game for T_{top} (T_4) and the resulting submatrix games for lower tiers. The minimax tier list is shown in Table 2, along with the associated minimax strategy for each tier. For this matrix, each tier has only one minimax strategy, so the ordering of agents in each tier corresponds to the ordering of weights (descending) in the minimax strategy shown.

The minimax strategy for the full game has a support containing seven agents. Having such a large top tier means that a majority of the agents should be seen used in optimal play, which can often be seen as a good thing. Just considering the remaining five agents, solving that matrix game produces a minimax strategy with the support $\{a_5, a_7, a_9\}$. Looking at the matrix M shown in Table 1, we can observe that these agents are able to share this tier because of their rock-paper-scissors relationship. Since there are only two unassigned agents remaining, the clear optimal strategy is to always select a_{10} because it beats a_6, that is, $M[10, 6] > 0$. As a result, the minimax tier list has two singleton tiers at its bottom.

Table 2. Resulting Minimax Tier List and Strategies

Tier	Agents	Minimax Strategy
T_4	$a_{11},\ a_8,\ a_1,\ a_3,$ $a_4,\ a_{12},\ a_2$	(.255, .247, .241, .078, .063, .062, .053)
T_3	$a_5,\ a_9,\ a_7$	(.667, .245, .088)
T_2	a_{10}	(1.0)
T_1	a_6	(1.0)

After generating 100 more matrices, we analyzed their resulting minimax tier lists. We observed an average of 4.39 tiers. T_{top} had an average size of 5.89, and $T_{\text{top}-1}$ had an average size of 2.96. The friendliness of the tier lists [2], defined as the number of non-singleton tiers divided by the total number of tiers, averaged 0.51. The minimax tier list shown in Table 2 is representative of our average results.

6 Conclusion

TCFGs provide a strong tier-based model for intransitive games. We have shown that previous TCFG preference frameworks have particularly relied on the assumption that agents have matchups in equal proportions, but this is often not the case. Our preference framework addresses these concerns by considering minimax strategies, which define optimal probabilities for players to select agents. We proved that our minimax preference framework is distinct from previous frameworks. We proved that each minimax tier list is the only Nash stable tier list and, furthermore, is the only core stable tier list for the provided payoff matrix. We gave an algorithm that generates this tier list in polynomial time, and we argued that the resulting tier lists can be meaningful in terms of optimal strategies for games.

Acknowledgments

We want to thank the reviewers for their comments and suggestions for improving the overall clarity of the paper.

References

[1] N. Arnold. *Tiered Coalition Formation Game Variants, Stability, and Simulation*. PhD thesis, University of Kentucky, 2024.
[2] N. Arnold, J. Goldsmith, and S. Snider. Extensions to tiered coalition formation games. In *The International FLAIRS Conference Proceedings*, volume 35, 2022.
[3] N. Arnold, N. Fluty, R. D. Flores, J. Goldsmith, and B. Harrison. Evaluating the stability of the smogon tier list for competitive pokémon battling, 2024. ExAG Workshop at AIIDE '24.
[4] N. Arnold, S. Snider, and J. Goldsmith. Socially conscious stability for tiered coalition formation games. *Annals of Mathematics and Artificial Intelligence*, 92(3):539–580, 2024.
[5] A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
[6] M. Bullinger. *Computing Desirable Outcomes in Coalition Formation*. PhD thesis, Technische Universität München, 2023.

[7] G. Chalkiadakis, E. Elkind, E. Markakis, and N. R. Jennings. Overlapping coalition formation. In *Internet and Network Economics: 4th International Workshop, WINE 2008, Shanghai, China, December 17-20, 2008. Proceedings 4*, pages 307–321. Springer, 2008.

[8] J. H. Dreze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society*, pages 987–1003, 1980.

[9] E. Elkind and J. Rothe. Cooperative game theory. *Economics and Computation: An introduction to algorithmic game theory, computational social choice, and fair division*, pages 135–193, 2016.

[10] D. Gale. *The theory of linear economic models*. University of Chicago press, 1989.

[11] P. Garner, A. Lin, R. Harris, B. Harrison, and J. Goldsmith. Arma 3 billeting as a coalition formation game, 2024. ExAG Workshop at AIIDE '24.

[12] J. Hajduková. Coalition formation games: A survey. *International Game Theory Review*, 8(04):613–641, 2006.

[13] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM Symposium on Theory of Computing*, pages 302–311, 1984.

[14] J. Schlueter, C. Addington, and J. Goldsmith. Anchored team formation games. In *The International FLAIRS Conference Proceedings*, volume 34, 2021.

[15] C. Siler. Tiered coalition formation games. In *FLAIRS*, pages 210–214, 2017.

[16] M. Spradling, J. Goldsmith, X. Liu, C. Dadi, and Z. Li. Roles and teams hedonic game. In *Algorithmic Decision Theory: Third International Conference, ADT 2013, Bruxelles, Belgium, November 12-14, 2013, Proceedings 3*, pages 351–362. Springer, 2013.

[17] J. Von Neumann and O. Morgenstern. Theory of games and economic behavior: 60th anniversary commemorative edition. In *Theory of Games and Economic Behavior*. Princeton university press, 2007.