Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)

Complexity of Hedonic Games with Dichotomous Preferences

Dominik Peters
Department of Computer Science
University of Oxford, UK
dominik.peters@cs.ox.ac.uk

Abstract

Hedonic games provide a model of coalition formation in which a set of agents is partitioned into coalitions and the agents have preferences over which set they belong to. Recently, Aziz et. al. (2014) have initiated the study of hedonic games with dichotomous preferences, where each agent either approves or disapproves of a given coalition. In this work, we study the computational complexity of questions related to finding optimal and stable partitions in dichotomous hedonic games under various ways of restricting and representing the collection of approved coalitions. Encouragingly, many of these problems turn out to be polynomial-time solvable. In particular, we show that an individually stable outcome always exists and can be found in polynomial time. We also provide efficient algorithms for cases in which agents approve only few coalitions, in which they only approve intervals, and in which they only approve sets of size 2 (the roommates case). These algorithms are complemented by NP-hardness results, especially for representations that are very expressive, such as in the case when agents' goals are given by propositional formulas.

Introduction

A *coalition* is an alliance between a group of individuals, formed in order to achieve a common goal. How do such coalitions form if agents are selfish? An extensive literature in economics and computer science has studied this question using the natural model of a *hedonic game* (see the survey byAziz and Savani (2016)). A hedonic game consists of a set of agents, each of which submits a preference ordering over all possible coalitions this agent could join. An outcome of the game is a partition of the agent set into disjoint coalitions. If agents are selfish, we want to find a *stable* outcome, while in other situations a welfare-optimal or fair outcome might be desired.

There have turned out to be multiple obstacles to achieving these tasks. First, not all hedonic games admit any stable outcome, and thus the search for one may be futile. Second, the computational problem of finding a partition that is stable, optimal, or fair has turned out to be intractable even for a large variety of severely restricted preference structures.

	SW	PF	PO	NS	IS	CR	SCR
Boolean	NP-c.	NP-c.	NP-h.	NP-c.	P	FNP-h.	Σ_2^p-c.*
1-lists	NP-c.	P	P	P	P	P	NP-c.
2-lists	NP-c.	P	P	P	P	P	NP-c.
3-lists	NP-c.	NP-c.	NP-h.	?	P	P	NP-c.
4-lists	NP-c.	NP-c.	NP-h.	NP-c.	P	P	NP-c.
Anonymous	NP-c.	NP-c.*	NP-h.	NP-c.	P	P	NP-c.
Intervals	P	P	P	?	P	P	?
Roommates	P	P	P*	NP-c.	P	P	P*
Majority	?	P	?	P	P	P	P

Table 1: Overview of complexity results for various dichotomous preference representations; results marked (*) were obtained elsewhere. The columns describe the problems of maximising welfare, and of finding (respectively) perfect, pareto-optimal, Nash-stable, individually stable, core-stable, and strict-core-stable partitions.

Recently, Peters and Elkind (2015) have shown that deciding whether a given hedonic game admits any stable outcome at all is NP-hard for preference restrictions and representations that allow agents to express more than 4 or 5 preference 'intensities' (with some mild additional qualifiers). A result of Deineko and Woeginger (2013) shows this to also be the case for a specific restriction allowing 3 intensities. These results suggest that if we want to stand a chance of finding polynomial time algorithms for a restricted class of hedonic games, we will need to go all the way down to dichotomous preferences, which allow only 2 preference intensities.

In the context of hedonic games, studying the restriction to dichotomous preferences has recently been proposed by Aziz et al. (2014). They represent agents' preferences by formulas of propositional logic. In particular, we can use the names of agents as propositional variables. An agent then approves a coalition if the members of that coalition satisfy her goal formula. Accordingly, they term games with this preference representation "boolean hedonic games".

This logic representation is attractive in that it is universally expressive for dichotomous preferences and often succinct. A further advantage is that we may use it to translate computational questions such as "find a stable partition" into propositional logic, and then use an off-the-shelf SAT solver to answer it. Such a translation is presented in detail by Aziz et al. (2014). Note that for some solution concepts, they have not found polynomial-size expressions. Still, given the impressive performance of modern SAT solvers on instances arising in practice, we might hope that this approach allows hedonic games to be applicable in practice.

Aziz et al. (2014) argue that their specific SAT encoding cannot be improved by an efficient algorithm if and only if the corresponding computational problems are NP-hard. They write that "identifying the complexity of finding partitions satisfying solution concepts for Boolean hedonic games is therefore the most immediate direction of further research."

In this work, we study a variety of different restrictions of agents' dichotomous preferences, and use the term "dichotomous hedonic game" for any hedonic game in which all agents have dichotomous preferences.

Theoretically speaking, the dichotomous case is nice because every dichotomous hedonic game admits an outcome that is simultaneously core-stable (resistant to group deviations) and individually stable (resistant to deviations by any single player). This is a refinement of an observation by Aziz et al. (2014). While the argument establishing the existence of such a partition does not yield a polynomial-time algorithm, we identify many appealing special cases in which there is one. In the general case, we always can in polynomial time find a partition that is individually stable (but not necessarily core-stable) when given an oracle that decides whether a given coalition is approved by a given agent.

We further study the computational complexity of finding a partition that maximises the number of players who approve it, of finding a Pareto optimal partition, and of deciding whether a strict-core-stable or a Nash-stable partition exists (these concepts are strengthenings of core- and individual stability). We find that these problems are all NP-hard under the logic representation (which is not too surprising given its expressive power) and welfare-maximisation in particular is inapproximable and fixed-parameter intractable. In contrast, we find that these problems become easy when requiring the collection of approved sets to have additional combinatorial structure. In particular, many of these problems are easy when agents are placed on a line and only approve intervals; when agents only approve sets of size at most 2 (the roommates case); when agents form a graph and approve sets in which they are connected to the majority of other vertices; and if agents only approve at most 2 different coalitions. That said, we also find isolated hardness results for these restrictions, and it turns out that requiring preferences to be *anonymous* (so that agents merely approve certain cardinalities of coalitions) does not make any of our problems easier.

These results, which are summarised in Table 1, establish dichotomous hedonic games as one of the very few subclasses of hedonic games that admit polynomial time solutions.

Preliminaries

A *hedonic game* $\langle N, (\succcurlyeq_i)_{i\in N}\rangle$ is given by a finite set N of agents, and for each agent $i \in N$ a complete and transitive preference relation over $\mathcal{N}_i = \{S \subseteq N : i \in S\}$. We write $\succ_i$ and $\sim_i$ for the strict and indifference parts of $\succcurlyeq_i$. A hedonic game has *dichotomous preferences*, and is called a *dichotomous hedonic game*, if for each agent $i \in N$, the coalitions $\mathcal{N}_i = \{S \subseteq N : i \in S\}$ can be partitioned into *approved* coalitions $\mathcal{N}_i^+$ and non-approved coalitions $\mathcal{N}_i^-$ such that i strictly prefers approved coalitions to non-approved coalitions, but is indifferent within the two groups: so $S \succ_i T$ iff $S \in \mathcal{N}_i^+$ and $T \in \mathcal{N}_i^-$.

The outcome of a hedonic game is a *partition* π of the agent set into disjoint coalitions. We write $\pi(i)$ for the coalition $S \in \pi$ that contains $i \in N$. We are interested in finding partitions that are *stable*, *optimal*, and/or *fair*. In a dichotomous hedonic game, a partition π *maximises social welfare* if it has the maximum number of agents who are in approved coalitions among all partitions of N. If *every* agent is in an approved coalition in π, then π is called *perfect* (sometimes known as *wonderfully stable*). A partition π is *Pareto-optimal* if there is no partition π' such that $\pi'(i) \succcurlyeq_i \pi(i)$ for all $i \in N$ and $\pi'(i) \succ_i \pi(i)$ for some $i \in N$. Fairness can be formalised using the notion of *envy-freeness*: a partition π is envy-free if there is no agent i who would prefer to be in the position of agent j, i.e. $\pi(j) \setminus \{j\} \cup \{i\} \succ_i \pi(i)$.

There are many notions of stability for a partition π in a hedonic game. We will mainly use four such concepts. A partition π is *core-stable* if there is no non-empty coalition $S \subseteq N$ with $S \succ_i \pi(i)$ for all $i \in S$. Thus, every member of S would strictly prefer being in S to being where they have been put under π. A partition π is *strict-core-stable* if there is no non-empty coalition $S \subseteq N$ with $S \succcurlyeq_i \pi(i)$ for all $i \in S$ and $S \succ_i \pi(i)$ for some $i \in S$. In both of these notions, a group of agent deviates. If we restrict our attention to the possibility of just a single agent deviating, we obtain the notion of *Nash-stability*. Here, no agent i prefers to join another coalition in π, that is $\pi(i) \succcurlyeq_i \pi(j) \cup \{i\}$ for all j, and also $\pi(i) \succcurlyeq_i \{i\}$. In an *individually stable (IS)* partition, no agent prefers to deviate in this way and is welcomed by his new coalition. Formally, an agent i *IS-deviates* into a coalition $S \in \pi \cup \{\varnothing\}$ if $S \succ_i \pi(i)$ and for each $j \in S$, we have $\pi(j) \cup \{i\} \succcurlyeq_j \pi(j)$. A partition is individually stable if no agent IS-deviates. If $\pi(i) \succcurlyeq_i \{i\}$ for all i, we also say that π is *individually rational*.

We will occasionally use the *Preference Refinement Algorithm* (PRA) for finding Pareto-optimal partitions in hedonic games, due to Aziz, Brandt, and Harrenstein (2013). Roughly, this algorithm is applicable to classes of hedonic games that are closed under refining agents' preferences $\succcurlyeq_i$ and further admit a polynomial-time algorithm that decides the existence of a perfect partition. Of course, dichotomous games are not closed under refining preferences (which introduces strict preferences), so PRA is only applicable if we can extend results to more general classes of hedonic games. A general observation from that paper is that for a class of games in which deciding the existence of a perfect partition is NP-hard, there is also no polynomial time algorithm that finds a Pareto-optimal partition (unless P = NP). This is the reason for the hardness results in the PO column of Table 1.

Existence Guarantees

Every dichotomous hedonic game admits a core-stable outcome (Aziz et al. 2014, Prop. 8). As we now show, we may additionally demand the outcome to be individually stable.

Proposition 1. *Every dichotomous hedonic game admits a partition that is both core-stable and individually stable.*

Proof. Repeatedly find a maximal *unanimous* coalition S (meaning it is approved by all its members), add S to π, and remove the agents in S from consideration. (S may be a singleton.) Once this is not possible anymore, put all remaining agents into a single coalition in π.

π is core-stable, since every possible blocking coalition S must consist entirely of dissatisfied players that all approve S. But then the procedure above wouldn't have put the agents in S into the losing coalition (because it is unanimous).

π is individually stable, since no dissatisfied agent is allowed to join any other coalition by maximality of the S selected, and no dissatisfied agent wants to defect into a singleton, since then the singleton coalition would have been assigned by the procedure. □

The partition produced in this proof is actually *strongly individually stable* as defined by Aziz and Brandl (2012), a concept which is stronger than both individual stability and core-stability.

Notice that the proof does not provide an algorithm for efficiently finding such a stable partition, since in general it is hard to decide whether there is still a set S left that is approved by each of its members. On the other hand, if the number of coalitions an agent approves is polynomially bounded, or if the cardinalities of approved coalitions are bounded by a constant, then the above procedure can be run in polynomial time.

The procedure given does provide an upper bound on the complexity of finding a core-stable outcome whenever we can decide in polynomial time whether a given agent approves a given coalition: then the problem is contained in the class FP^{NP}, the class of function problems solvable by a polynomial-time algorithm when given an NP-oracle. The NP-oracle in this case is used to return a unanimous coalition or report that there is none.

If we drop the requirement of core-stability, we obtain a more encouraging result. As before, we assume that we can decide in polynomial time whether a given agent approves a given coalition. All reasonable representations of dichotomous hedonic games—and all representations we consider—will have this property.

Proposition 2. *For every dichotomous hedonic game, we can find an individually stable partition in $O(n^3)$ calls to an oracle that decides whether a given coalition $S \subseteq N$ is approved by a given agent $i \in N$.*

Proof. Run Algorithm 1 which simulates successive deviations by single agents. Note that the **while** loop executes at most n times, since each agent is assigned only once. □

Algorithm 1 Find an individually stable partition

$\pi \leftarrow \varnothing$
set every agent to unassigned
for each agent i that approves $\{i\}$ **do**
  assign i to coalition $\{i\}$
while there is i unassigned
    who can IS-deviate into $\pi(j)$ **do**
  assign i to coalition $\pi(j)$
assign all unassigned agents into a single coalition
return π.

Boolean Hedonic Games

Aziz et al. (2014) study a representation of dichotomous hedonic games using propositional formulas that is both universally expressive and often succinct when the collections $\mathcal{N}_i^+$ of approved sets have a combinatorial structure. The logic representation allows each agent i to submit a logical formula ϕ_i—agent i's *goal*—such that the agent approves a coalition S if and only if ϕ_i evaluates to true on S. More specifically, we use propositional logic where the set of propositional atoms[1] (i.e., variables) is given by the agent set N. A formula ϕ of this logic is *satisfied* by $S \subseteq N$, written $S \models \phi$, if and only if the formula is true under the assignment that sets variables i with $i \in S$ to true and variables j with $j \notin S$ to false. If agent i has the goal ϕ_i, then i's approved coalitions are those that satisfy i's goal:

$$S \in \mathcal{N}_i^+ \iff S \models \phi_i.$$

For example, if $\phi_1 = 2 \vee (\neg 2 \wedge 3) \vee 4$, then agent 1 approves coalitions as soon as they contain either agent 2 or 4, but is also happy when agent 3 is present but 2 is not.

With this definition, a dichotomous hedonic game can be specified by a list of formulas, one formula per agent, and is then called a *boolean hedonic game*. Note that by using formulas in disjunctive normal form (DNF), this representation is universally expressive.

Given the use of logic, it is perhaps unsurprising that many computational problems about boolean hedonic games in logic representation are computationally hard. We now give a selection of results to this effect.

Recall that Proposition 1 shows that every dichotomous hedonic game admits a core-stable partition. It turns out that, despite this fact, for boolean hedonic games we cannot find a core-stable partition in polynomial time unless P = NP.

Theorem 3. *It is FNP-hard to find a core-stable partition in a boolean hedonic game.*

Proof. We reduce from FSAT, the function problem of finding a satisfying assignment for a given propositional formula. So let ϕ be a formula with n variables. For each variable x occurring in ϕ, introduce an agent x. Every agent in this game has as her goal the formula ϕ. Because every agent has the same goal, every coalition $S \subseteq N$ is unanimously approved

[1] Aziz et al. (2014) use variables of form $\{i, j\}$ which is more convenient for logical characterisations.

or disapproved. We need to show how to transform a core-stable outcome of this game into a satisfying assignment of ϕ or into a correct report that there is none.

Let π be a core-stable partition. Go through the at most n coalitions in π, and try to find a coalition $S \in \pi$ that is approved by all its members. If this exists, then the assignment that sets variables in S to true and variables outside S to false satisfies ϕ. If no such coalition exists, then ϕ is unsatisfiable, because if there was a satisfying assignment α, then the variables true in α would form a unanimously approved coalition which would block π. ☐

Corollary 4. *It is coNP-complete to decide whether a given partition π is core-stable in a boolean hedonic game.*

Proof. In the reduction above, the all-singletons partition is blocked only by satisfying assignments of ϕ (we may assume that no singleton assignment satisfies ϕ). ☐

We have shown that finding a core-stable partition is FNP-hard, and is contained in FP^{NP}. Because of Corollary 4, the problem is unlikely to be contained in FNP. We leave open the problem of pinpointing the complexity of this problem.

Peters (2015) shows that deciding the existence of a strict-core-stable partition is Σ_2^p-complete for boolean hedonic games. If we are only interested in NP-hardness, reductions like the one above can be adapted for the strict-core. For other solution concepts, hardness for the logic representation follows by generalisation of problems proven hard for more restricted preference classes below.

Lists

In a context in which it is sensible to presume that agents will only approve at most polynomially many coalitions, we can represent their preferences by merely listing all approved coalitions. The complexity of stability problems for lists in the non-dichotomous case is studied by Ballester (2004). We consider here an even more restricted variant: in the k-list representation, every agent submits a list of at most k approved coalitions. If $k \leqslant \ell$, notice that a hardness result for k-lists also applies to ℓ-lists, and that poly-time algorithms for ℓ-lists also work for k-lists. As observed above, finding a core-stable partition is easy for k-lists (or even poly(n)-lists).

Perhaps surprisingly, we already have some hardness results in the case where agents approve only a single coalition.

Theorem 5. *Maximising social welfare is NP-complete even for 1-lists.*

Proof. Reduce from INDEPENDENT SET. Given a graph $G = (V, E)$ and target size k, we produce a game with 1-lists that admits an outcome with $\geqslant k$ satisfied agents if and only if G contains an independent set of size $\geqslant k$. We introduce one agent for each vertex and one agent for each edge. We list the edges as $e_1, \dots, e_m$. The edge agents e_i submit empty lists: they do not approve any coalition. The vertex agent v approves $A_v := \{v\} \cup \{e_i \in E : v \in e_i\}$, that is, v approves being together with the edges incident to it.

Suppose G contains the independent set $U \subseteq V$ with $|U| \geqslant k$. Then take the partition π consisting of coalitions A_v for each $v \in U$, and singleton coalitions for everyone not in $\bigcup_{v \in U} A_v$. The sets listed are disjoint because U is independent. Clearly, π delivers welfare k.

Suppose the game admits a partition π with welfare at least k. Thus at least k vertex agents are in an approved coalition in π. Since the approved sets of adjacent vertices intersect (in the edge agent connecting them), this means that the vertices in approved sets form an independent set of size at least k. ☐

From known inapproximability results for MAX-INDEPENDENT-SET, we may deduce that maximising social welfare is not approximable within $n^{1/2-\varepsilon}$ for any $\varepsilon > 0$ unless NP = ZPP (Håstad 1999). In the other direction, we *can* approximate the problem using approximation algorithms for WEIGHTED SET-PACKING, where coalitions are weighted by the number of agents approving it. This gives a $\sqrt{n}$-approximation (Halldórsson 2000), which works whenever lists have polynomial length.

Since INDEPENDENT SET is $W[1]$-hard (Downey and Fellows 1995), the reduction also shows that maximising welfare is $W[1]$-hard with parameter the number of approving agents.

Theorem 6. *Deciding existence of a strict-core-stable partition is NP-complete even for 1-lists.*

Proof. The proof is similar to the previous one, replacing INDEPENDENT SET by the decision problem KERNEL, with arc agents (u, v) approving A_u. We omit the details due to space restrictions. ☐

Finding partitions satisfying other concepts only becomes hard for 3-lists, while 2-lists admit efficient algorithms.

Theorem 7. *Finding a perfect partition or a Nash-stable partition is easy for 2-lists.*

Proof. We reduce to 2SAT. For perfect partitions, let $\mathcal{A}$ be the collection of all coalitions that appear on the lists. For all pairs $S, T \in \mathcal{A}$ of intersecting partitions, add a clause $(\neg S \vee \neg T)$. For each agent i approving A and B, add a clause $(A \vee B)$. Then any satisfying assignment can be translated to a perfect partition.

For Nash-stability, let $\mathcal{B}$ be the collection of all approved sets, and also of sets $S_{-i} := S \setminus \{i\}$ for coalitions S approved by i. As before, add clauses $(\neg S \vee \neg T)$ for intersecting coalitions in $\mathcal{B}$. But now, for each agent i approving sets A and B, add clauses $(A_{-i} \to B)$ and $(B_{-i} \to A)$. This expresses Nash-stability. ☐

We can use a similar technique to express envy-freeness in 2SAT. Thus, in the 2-list case, we can also look for perfect or Nash-stable partitions that are additionally envy-free.

For perfect partitions, the reduction to 2SAT works even for preferences starting in $S \succ_i T \succ_i \{i\}$, i.e., even after refining a 2-list. This means that we can use the preference refinement algorithm (PRA) to find a Pareto-optimal partition for games given by 2-lists.

As might be expected due to our use of 2SAT, these easiness results do not extend to $k = 3$.

Theorem 8. *Deciding existence of a perfect partition or a strict-core-stable partition is NP-complete for 3-lists.*

Proof. A given partition can be checked to be perfect by verifying that every agent is in a coalition appearing in his 3-list; it can be checked to be strict-core-stable by making sure that no coalition that appears in a list weakly blocks. NP-hardness follows by a reduction from X3C restricted to each element appearing in at most 3 sets. Given elements $X = \{x_1, \dots, x_{3n}\}$ and sets $S = \{s_1, \dots, s_m\}$, take the game with agent set X, with each agent $x_i \in X$ approving of the at most 3 sets from S it appears in. Notice that any strict-core-stable partition in this game must also be perfect, for if any element x_i is not satisfied then it can weakly block with a set containing it. Thus, a partition of the agent set is strict-core-stable iff it is perfect iff it is a solution of the X3C-instance. $\square$

Deciding the existence of a Nash-stable partition is NP-complete for 4-lists. This follows from the result for the roommate case in Theorem 11. We were unable to decide the complexity of Nash-stability in the 3-list case.

Anonymous Preferences

In an *anonymous* hedonic game, agents' preferences $\succcurlyeq_i$ are determined by an underlying ordering $\trianglerighteq_i$ over the possible coalition sizes $\{1, \dots, |N|\}$, with $S \succcurlyeq_i T$ iff $|S| \trianglerighteq_i |T|$. Ballester (2004) has shown that the problems of deciding existence of core-, Nash-, and individually stable partitions are all NP-complete for anonymous preferences. We show hardness of the existence of perfect, strict-core- and Nash stable partitions even in the case when $\succcurlyeq_i$ (and thus $\trianglerighteq_i$) are dichotomous. By contrast, it is easy to find an individually stable partition (as always), and in the anonymous case it is also easy to decide whether a unanimous coalition exists (by counting the number of agents that approve each coalition size), so that we can find a core-stable partition in poly time.

Theorem 9. *The problems of deciding existence of a perfect, a strict-core-stable, or a Nash-stable partition are NP-complete for anonymous preferences, even if at most 4 sizes are approved.*

Proof. For perfect partitions, this result is stated without proof as Theorem 4.4 by Darmann et al. (2012). A proof appears in the survey by Woeginger (2013). For the strict core, the same reduction can be used.

For Nash-stability, we reduce from X3C. Given elements $X = \{x_1, \dots, x_{3n}\}$ and sets $S = \{s_1, \dots, s_m\}$, assign the *code number* $12k$ to the set s_k. Introduce $3n$ agents, one for each element; for each set s_k introduce $12k - 3$ dummy agents; finally introduce 1 stalker player. The element player of x_i approves coalition sizes 1 and the code numbers of the sets that include x_i. The dummy agents of s_k approve sizes 1, 3, and $12k$. The stalker agent approves size 2.

Suppose the X3C-instance admits a solution. For each set s_k that is used in the solution, make a coalition consisting of the 3 element players appearing in s_k and the $12k - 3$ dummies of s_k, making a coalition of $12k$ happy players. For each set s_k that is not used in the solution, arbitrarily match its $12k - 3$ dummies into triples. The stalker forms a singleton coalition. Then everyone but the stalker is happy, and since no other coalition is of size 1, the stalker does not Nash deviate. Hence the resulting partition is Nash-stable.

Conversely suppose there is a Nash-stable partition π. Since both elements and dummies approve coalition size 1, they must be in a coalition of an approved size in π. If any of them were in a singleton, then the stalker player would join them; hence they are not. It follows that the elements and dummy players are all in coalitions of size a multiple of 3, and thus, since the number of element and dummy players together is also a multiple of 3, the stalker must be in a singleton coalition in π. But now we can extract a solution for the X3C-instance from π by taking for each element its coalition size, which codes for the set it is in. $\square$

Intervals

Suppose that the agent set can be put in some linear order, say $N = \{1, 2, \dots, n\}$ with the natural ordering. Suppose further that each agent i only approves intervals $[a, b]$ of agents (with $a \leqslant i \leqslant b$). In a restriction like this, termed "candidate interval (CI)" by Elkind and Lackner (2015) and also studied by Faliszewski et al. (2009), dynamic programming promises to be of help, and indeed this is the case. But note first that since there are only $O(n^2)$ possible approved coalitions (choose $\binom{n}{2}$ interval endpoints), we can quickly find a core-stable outcome in this restriction.

Theorem 10. *If every agent only approves intervals, we can in polynomial time find a welfare-maximising partition.*

Proof. We give a dynamic programming algorithm. For each $m = 0, \dots, n$, let $W[m]$ denote the maximum welfare obtainable in the subgame obtained by restricting to the agent set $\{1, \dots, m\}$, with each agent approving all originally approved coalitions S such that $S \subseteq \{0, \dots, m\}$. In particular we have $W[0] = 0$ (the empty sum), and $W[n]$ is the sought value. Write $\#[s, m]$ for the number of agents that approve the interval $[s, m]$ in this subgame. We then have the following recurrence:

$$W[m] = \max_{s=1,\dots,m} \Big\{ W[s-1] + \#[s, m] \Big\}.$$

To see this, note that there always is a welfare-maximising partition in which all coalitions are intervals (since we can weakly increase the welfare by splitting any coalition into its 'connected components'). In this partition, m must be part of some interval $[s, m]$, and the subgame $\{1, \dots, s-1\}$ must be in welfare maximum. $\square$

Since we can quickly maximise welfare, we can also quickly decide the existence of a perfect partition. Notice that the algorithm presented can easily be adapted to work even in the non-dichotomous case (by replacing $\#[s, m]$ by the total utility accrued in this coalition), and a further slight modification allows us to decide existence of a perfect partition even in this general case. Using this, it follows that by using the PRA algorithm, we can in polynomial time find a Pareto-optimal partition when agents only approve intervals (this result also carries over to the non-dichotomous case). We leave open the complexity of deciding the existence of Nash-stable and strict-core-stable partitions in this case, both of which seem resistant to naïve dynamic programming and greedy algorithms. For a related setting in which no disconnected coalitions are allowed, Igarashi and Elkind (2016) have a positive result.

Roommates

Let us now consider the restriction of dichotomous hedonic games where agents only approve coalitions of size at most 2. This case could also be referred to as a (stable) roommate problem with dichotomous preferences (a special case of SRTI). By the analysis of Scott (2005), a strict-core-stable partition may be found (if it exists) in $O(m^2)$ time where m is the total number of approved pairs. If we additionally require preferences to be *mutual*, so that no pair $\{i, j\}$ is ever approved by only 1 partner, we can improve this runtime to $O(m\sqrt{n})$ (Abraham et al. 2007). By finding maximum weighted matchings, we can in $O(m\sqrt{n})$ time find a maximum welfare partition (and thus decide the existence of a perfect partition). Using the algorithms in Aziz, Brandt, and Harrenstein (2013), we can also find a Pareto-optimal partition. All these algorithms work even in the non-dichotomous case. Restricted to the dichotomous case, we can of course find individually stable and core-stable partitions, since at most n^2 coalitions are approved. Despite all this good news, there is also the following.

Theorem 11. *Deciding the existence of a Nash stable partition for dichotomous roommates is NP-complete, even in the bipartite (marriage) case, and even if each agent approves at most 4 coalitions.*

Proof. The proof is a dichotomization of the reduction in Theorem 2 of Peters and Elkind (2015); we omit the details here. For the non-dichotomous case, a similar result is given by Aziz (2013). □

Majority Games

In this section, we consider agents that approve those coalitions in which the agent is friends with the majority of members. More precisely, suppose we are given a graph $G = (N, E)$, where each agent corresponds to a vertex. Two agents are (mutually) friends if there is an edge between them. An agent i approves the coalition $S \ni i$ if $d_i(S) \geqslant |S|/2$, that is if i is connected to at least $|S|/2$ of the vertices in S. We refer to such a game as a *majority game*. This setting appears rather natural: in a group where the majority of agents are similar to me, I can hope that the group will agree and vote with me on issues that need to be decided. This class can also be seen as a dichotomisation of 'fractional hedonic games' (Aziz, Brandt, and Harrenstein 2014).

Theorem 12. *In a majority game, a partition that is both Nash-stable and core-stable is guaranteed to exist and can be found in polynomial time.*

Proof. Take π be a maximum-cardinality packing in G of vertex-disjoint edges and triangles. Such a maximum packing can be found by an alternating-path-algorithm due to Hell and Kirkpatrick (1984). The set of agents not put into an edge or triangle forms an independent set of G, so that there cannot be a core-deviation from within this set. Also, there cannot be Nash deviations by maximality of the packing. □

The partition produced in the proof above is actually even *strongly Nash stable* in the sense of Karakaya (2011).

It turns out that the same algorithm can be used to decide the existence of a perfect and of a strict-core-stable partition. To see this, we appeal to two famous results from extremal graph theory. Let $\delta(G)$ and $\Delta(G)$ be the minimum and maximum degree of G, with n the number of vertices of G.

Theorem (Dirac)**.** *If $\delta(G) \geqslant \frac{n}{2}$ then G is Hamiltonian.*

Theorem (Hajnal-Szemerédi)**.** *If $\Delta(G) \leqslant k - 1$ then G has an equitable k-colouring.*

An *equitable k-colouring* is a proper vertex colouring of G using k colours and such that the size of any two colour classes differs by at most 1. The Hajnal-Szemerédi theorem was conjectured by Paul Erdős and proved in (Hajnal and Szemerédi 1970). These two theorems imply the following.

Lemma 13. *If G is a graph with $\delta(G) \geqslant \frac{n}{2}$ then the vertices of G can be partitioned into edges and triangles.*

Proof. If n is even, then take every other edge on a Hamilton cycle of G which is a perfect matching. If n is odd, say $n = 2k + 1$, then consider the complement graph $\overline{G}$ and note

$$\Delta(\overline{G}) = |G| - \delta(G) - 1 \leqslant n - \tfrac{n}{2} - 1 = k - \tfrac{1}{2}.$$

Since $\Delta(\overline{G})$ is integer, it follows that $\Delta(\overline{G}) \leqslant k - 1$. By the Hajnal-Szemerédi theorem, $\overline{G}$ has an equitable k-colouring which is a partition of $\overline{G}$ into k independent sets $\overline{K}_{s_i}$ of sizes $s_1, \dots, s_k$. Taking complements, G can be partitioned into k cliques K_{s_i}. Since this partition is equitable, $|s_i - s_j| \leqslant 1$ for all $1 \leqslant i, j \leqslant k$. Since $\sum_{i=1}^{k} s_i = n = 2k + 1$, it follows that all $s_i \in \{2, 3\}$, as required. □

This lemma implies that without loss of generality a perfect partition consists of edges and triangles.

Theorem 14. *In majority games, perfect and strict-core-stable partitions coincide. If a perfect partition exists, then a perfect partition consisting of edges and triangles exists. Hence there is a polynomial time algorithm which will produce a perfect and strict-core-stable partition if it exists.*

Proof. By considering connected components of G separately, we may assume that G is connected. We may also assume that G does not contain isolated vertices. Then if π were strict-core-stable but not perfect, an edge with one unsatisfied endpoint would weakly block, a contradiction. If π is perfect, then the induced subgraphs $G[S]$ for each $S \in \pi$ satisfy the condition of Lemma 13, and thus S can be partitioned into edges and triangles preserving perfection.

Thus, a perfect and strict-core-stable partition exists if and only if a perfect packing of edges and triangles exists which can be checked using Hell and Kirkpatrick's algorithm. □

Conclusions

We have investigated a variety of preference representations and restrictions for hedonic games with dichotomous preferences, and have uncovered quite an interesting complexity landscape. Embedded in the wider hedonic games literature, this landscape is unusually encouraging, with many polynomial time solutions in interesting subclasses available. Still, we have to leave a number of interesting questions open; in particular we were unable to settle the 'frontier of tractability' for Nash-stability in k-lists, and we do not know how to decide the existence of Nash and strict-core-stable partitions in the interval case. A further avenue for future work is defining additional appealing preference restrictions and studying their complexity properties.

Taking a broader view, it would be interesting to see an empirical evaluation of the power of boolean hedonic games when solved by modern SAT solvers. Besides implementation, a major challenge here seems to be obtaining sensible, preferably even real-world, preference data for this domain.

Acknowledgements Dominik Peters is supported by EPSRC, and received additional support through COST Action IC1205 on Computational Social Choice. The author is grateful for helpful suggestions by the AAAI reviewers, and thanks Edith Elkind, Paul Harrenstein, and Ayumi Igarashi for useful discussions.

References

Abraham, D. J.; Levavi, A.; Manlove, D. F.; and OMalley, G. 2007. The stable roommates problem with globally-ranked pairs. In *Internet and Network Economics*. Springer. 431–444.

Aziz, H., and Brandl, F. 2012. Existence of stability in hedonic coalition formation games. AAMAS '12, 763–770.

Aziz, H., and Savani, R. 2016. Hedonic games. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. chapter 15.

Aziz, H.; Harrenstein, P.; Lang, J.; and Wooldridge, M. 2014. Boolean hedonic games. LOFT '14, arXiv:1509.07062 [cs.GT].

Aziz, H.; Brandt, F.; and Harrenstein, P. 2013. Pareto optimality in coalition formation. *Games and Economic Behavior* 82:562–581.

Aziz, H.; Brandt, F.; and Harrenstein, P. 2014. Fractional Hedonic Games. AAMAS '14, 5–12.

Aziz, H. 2013. Stable marriage and roommate problems with individual-based stability. AAMAS '13, 287–294.

Ballester, C. 2004. NP-completeness in hedonic games. *Games and Economic Behavior* 49(1):1–30.

Darmann, A.; Elkind, E.; Kurz, S.; Lang, J.; Schauer, J.; and Woeginger, G. 2012. Group activity selection problem. In *Internet and Network Economics*. Springer. 156–169.

Deineko, V. G., and Woeginger, G. J. 2013. Two hardness results for core stability in hedonic coalition formation games. *Discrete Applied Mathematics* 161(13):1837–1842.

Downey, R. G., and Fellows, M. R. 1995. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science* 141(1):109–131.

Elkind, E., and Lackner, M. 2015. Structure in dichotomous preferences. In *IJCAI '15*, 2019–2025.

Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2009. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge*, 118–127. ACM.

Hajnal, A., and Szemerédi, E. 1970. Proof of a conjecture of Erdős. *Combinatorial theory and its applications* 2:601–623.

Halldórsson, M. M. 2000. Approximations of weighted independent set and hereditary subset problems. *Journal of Graph Algorithms and Applications* 4(1):1–16.

Hell, P., and Kirkpatrick, D. G. 1984. Packings by cliques and by finite families of graphs. *Discrete Mathematics* 49(1):45–59.

Håstad, J. 1999. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica* 182(1):105–142.

Igarashi, A., and Elkind, E. 2016. Hedonic games with graph-restricted communication. (in preparation).

Karakaya, M. 2011. Hedonic coalition formation games: A new stability notion. *Mathematical Social Sciences* 61(3):157–165.

Peters, D., and Elkind, E. 2015. Simple Causes of Complexity in Hedonic Games. IJCAI '15, 617–623.

Peters, D. 2015. Σ_2^p-complete problems on hedonic games. arXiv:1509.02333 [cs.GT].

Scott, S. 2005. *A study of stable marriage problems with ties*. Ph.D. Dissertation, University of Glasgow.

Woeginger, G. J. 2013. Core stability in hedonic coalition formation. In *SOFSEM 2013: Theory and Practice of Computer Science*. Springer. 33–50.