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

Computing Possible and Necessary Equilibrium Actions (and Bipartisan Set Winners)

Markus Brill, Rupert Freeman and **Vincent Conitzer**
Department of Computer Science
Duke University
Durham, NC 27708, USA
{brill,rupert,conitzer}@cs.duke.edu

Abstract

In many multiagent environments, a designer has some, but limited control over the game being played. In this paper, we formalize this by considering incompletely specified games, in which some entries of the payoff matrices can be chosen from a specified set. We show that it is NP-hard for the designer to make this choices optimally, even in zero-sum games. In fact, it is already intractable to decide whether a given action is (potentially or necessarily) played in equilibrium. We also consider incompletely specified symmetric games in which all completions are required to be symmetric. Here, hardness holds even in weak tournament games (symmetric zero-sum games whose entries are all -1, 0, or 1) and in tournament games (symmetric zero-sum games whose non-diagonal entries are all -1 or 1). The latter result settles the complexity of the possible and necessary winner problems for a social-choice-theoretic solution concept known as the bipartisan set. We finally give a mixed-integer linear programming formulation for weak tournament games and evaluate it experimentally.

1 Introduction

Game theory provides the natural toolkit for reasoning about systems of multiple self-interested agents. In some cases, the game is exogenously determined and all that is left to do is to figure out how to play it. For example, if we are trying to solve heads-up limit Texas hold'em poker (as was recently effectively done by Bowling et al., 2015), there is no question about what the game is. Out in the world, however, the rules of the game are generally not set in stone. Often, there is an agent, to whom we will refer as the *designer* or *principal*, that has some control over the game played. Consider, for example, applications of game theory to security domains (Pita et al. 2009; Tsai et al. 2009; An et al. 2012; Yin et al. 2012). In the long run, the game could be changed, by adding or subtracting security resources (Bhattacharya, Conitzer, and Munagala 2011) or reorganizing the targets being defended (roads, flights, etc.).

Mechanism design constitutes the extreme case of this, where the designer typically has complete freedom in choosing the game to be played by the agents (but still faces a challenging problem due to the agents' private information). However, out in the world, we generally also do not find this other extreme. Usually, some existing systems are in place and place constraints on what the designer can do. This is true to some extent even in the contexts where mechanism design is most fruitfully applied. For example, one can imagine that it would be difficult and costly for a major search engine to entirely redesign its existing auction mechanism for allocating advertisement space, because of existing users' expectations, interfacing software, etc. But this does not mean that aspects of the game played by the advertisers cannot be tweaked in the designer's favor.

In this paper, we introduce a general framework for addressing intermediate cases, where the designer has some but not full control over the game. We focus on *incompletely specified games*, where some entries of the game matrix contain *sets* of payoffs, from among which the designer must choose. The designer's aim is to choose so that the resulting equilibrium of the game is desirable to her. This problem is conceptually related to k-implementation (Monderer and Tennenholtz 2004) and the closely related internal implementation (Anderson, Shoham, and Altman 2010), where one of the parties is also able to modify an existing game to achieve better equilibria for herself. However, in those papers the game is modified by committing to payments, whereas we focus on choosing from a fixed set of payoffs in an entry.

We focus on two-player zero-sum games, both symmetric and not (necessarily) symmetric, and show NP-hardness in both cases. (Due to a technical reason explained later, hardness for the symmetric case does not imply hardness for the not-necessarily-symmetric case.) The hardness result for the symmetric case holds even for *weak tournament games*, in which the payoffs are all -1, 0, or 1, and for *tournament games*, in which the off-diagonal payoffs are all -1 or 1.

These results have direct implications for related problem in *computational social choice*, another important research area in multiagent systems. In social choice (specifically, voting), we take as input a vector of rankings of the alternatives (e.g., $a \succ c \succ b$) and as output return some subset of the alternatives. Some social choice functions are based on the *pairwise majority graph* which has a directed edge from one alternative to another if a majority of voters prefers the former. One attractive concept is that of the *essential set* (Laffond, Laslier, and Le Breton 1993a;

Dutta and Laslier 1999), which can be thought of as based on the following weak tournament game. Two abstract players simultaneously pick an alternative, and if one player's chosen alternative has an edge to the other's, the former wins. Then, the essential set (ES) consists of all alternatives that are played with positive probability in some equilibrium. In the absence of majority ties, this game is a tournament game and its essential set is referred to as the *bipartisan set* (BP).

An important computational problem in social choice is the *possible (necessary) winner problem* (Conitzer and Sandholm 2002; Konczak and Lang 2005; Lang et al. 2012; Xia and Conitzer 2011; Aziz et al. 2012): given only *partial* information about the voters' preferences—for example, because we have yet to elicit the preferences of some of the voters—is a given alternative potentially (necessarily) one of the chosen ones? It can thus be seen that our hardness results for incompletely specified (weak) tournament games directly imply hardness for the possible/necessary winner problems for ES and BP.

We conclude the paper by formulating and evaluating the efficacy of a mixed-integer linear programming formulation for the possible equilibrium action problem in weak tournament games. Due to the space constraint, most (details of) proofs have been omitted and can be found in the full version of this paper.

2 Examples

The following is an incompletely specified two-player symmetric zero-sum game with actions a, b, c, d.

	a	b	c	d
a	0	1	0	$\{-1,0,1\}$
b	-1	0	1	0
c	0	-1	0	1
d	$\{-1,0,1\}$	0	-1	0

Here, each entry specifies the payoff to the row player (since the game is zero-sum, the column player's payoff is implicit), and the set notation indicates that the payoff in an entry is not yet fully specified. E.g., $\{-1, 0, 1\}$ indicates that the designer may choose either -1, 0, or 1 for this entry. In the case of symmetric games, we require that the designer keep the game symmetric, so that if she sets[1] $u_r(d, a) = 1$ then she must also set $u_r(a, d) = -1$. Thus, our example game has three possible completions. The goal for the designer, then, is to choose a completion in such a way that the equilibrium of the resulting game is desirable to her. For example, the designer may aim to have only actions a and c played with positive probability in equilibrium. Can she set the payoffs so that this happens? The answer is yes, because the completion with $u_r(a, d) = 1$ has this property. Indeed, for any $p \geq \frac{1}{2}$, the mixed strategy $pa+(1-p)c$ is an equilibrium strategy for this completion (and no other equilibrium strategies exist). On the other hand, the completion with $u_r(a, d) = -1$ does have Nash equilibria in which b and d are played with positive probability (for example, both players mixing uniformly is an equilibrium of this game).

Next, consider the following incompletely specified asymmetric zero-sum game:

	ℓ	r
t	-2	1
b	$\{-1,1\}$	0

Suppose the designer's goal is to *avoid* row t being played in equilibrium. One might think that the best way to achieve this is to make row b (the only other row) look as good as possible, and thus set $u_r(b, \ell) = 1$. This results in a fully mixed equilibrium where t is played with probability $\frac{1}{4}$ (and ℓ with $\frac{1}{4}$). On the other hand, setting $u_r(b, \ell) = -1$ results in ℓ being a strictly dominant strategy for the column player, and thus the row player would actually play b with probability 1.

3 Preliminaries

In this section, we formally introduce the concepts and computational problems studied in the paper. For a natural number n, let $[n]$ denote the set $\{1, \ldots, n\}$.

3.1 Games

A matrix $M \in \mathbb{Q}^{m\times n}$ defines a two-player zero-sum game (or *matrix game*) as follows. Let the rows of M be indexed by $I = [m]$ and the columns of M be indexed by $J = [n]$, so that $M = (m(i, j))_{i\in I,j\in J}$. Player 1, the *row player*, has action set I and player 2, the *column player*, has action set J. If the row player plays action $i \in I$ and the column player plays action $j \in J$, the payoff to the row player is given by $m(i, j)$ and the payoff to the column player is given by $-m(i, j)$. A *(mixed) strategy* of the row (resp., column) player is a probability distribution over I (resp., J). Payoffs are extended to mixed strategy profiles in the usual way.

A matrix game $M = (m(i, j))_{i\in I,j\in J}$ is *symmetric* if $I = J$ and $m(i, j) = -m(j, i)$ for all $(i, j) \in I \times J$. A *weak tournament game* is a symmetric matrix game in which all payoffs are from the set $\{-1, 0, 1\}$. Weak tournament games naturally correspond to directed graphs $W = (A, \succ)$ as follows: vertices correspond to actions and there is a directed edge from action a to action b (denoted $a \succ b$) if and only if the payoff to the row player in action profile (a, b) is 1. A *tournament game* is a weak tournament game with the additional property that the payoff is 0 *only if* both players choose the same action. The corresponding graph thus has a directed edge for every pair of (distinct) vertices.

3.2 Incomplete Games

An *incompletely specified matrix game* (short: *incomplete matrix game*) is given by a matrix $M \in (2^{\mathbb{Q}})^{m\times n}$. That is, every entry of the matrix $M = (m(i, j))_{i\in I,j\in J}$ is a subset $m(i, j) \subseteq \mathbb{Q}$. If $m(i, j)$ consists of a single element, we say that the payoff for action profile (i, j) is *specified*, and write $m(i, j) = m$ instead of the more cumbersome $m(i, j) = \{m\}$. For an incomplete matrix game, the set of *completions* is given by the set of all matrix games that arise

[1]Let $u_r(x, y)$ denote the payoff to the row player in row x and column y.

from selecting a number from the corresponding set for every action profile for which the payoff is unspecified.

An *incomplete symmetric game* is an incomplete matrix game with $m(j,i) = \{-m : m \in m(i,j)\}$ for all $i \in I$ and $j \in J$. The set of *symmetric completions* of an incomplete symmetric game is given by the set of all completions that are symmetric. When considering incomplete symmetric games, we will restrict attention to symmetric completions, which is the reason hardness results do not transfer from the symmetric case to the general case. An *incomplete weak tournament game* is an incomplete symmetric game for which (1) every unspecified payoff has the form $m(i,j) = \{-1, 0, 1\}$ with $i \neq j$, and (2) every symmetric completion is a weak tournament game. An *incomplete tournament game* is an incomplete symmetric game for which (1) every unspecified payoff has the form $m(i,j) = \{-1, 1\}$ with $i \neq j$, and (2) every symmetric completion is a tournament game. Every incomplete (weak) tournament game corresponds to a directed graph in which the relation for certain pairs (i,j) of distinct vertices is unspecified. Whereas every completion of an incomplete tournament game satisfies either $m(i,j) = 1$ or $m(i,j) = -1$ for any such pair, a completion of an incomplete *weak* tournament game also allows "ties," i.e., $m(i,j) = 0$.

3.3 Equilibrium Concepts

The standard solution concept for normal-form games is Nash equilibrium. A strategy profile (σ, τ) is a Nash equilibrium of a matrix game M if the strategies σ and τ are best responses to each other, i.e., $m(\sigma, j) \geq m(\sigma, \tau) \geq m(i, \tau)$ for all $i \in I$ and $j \in J$. The payoff to the row player is identical in all Nash equilibria, and is known as the *value* of the game.

We are interested in the question whether an action is played with positive probability in at least one Nash equilibrium. For improved readability, the following definitions are only formulated for the row player; definitions for the column player are analogous. The support $supp(\sigma)$ of a strategy σ is the set of actions that are played with positive probability in σ. An action $i \in I$ is called *essential* if there exists a Nash equilibrium (σ, t) with $i \in supp(\sigma)$. By $ES_{\text{row}}(M)$ we denote the set of all actions $i \in I$ that are essential.

Definition 1. *The* essential set $ES(M)$ *of a matrix game* M *contains all actions that are essential, i.e.,* $ES(M) = ES_{row}(M) \cup ES_{column}(M)$.

There is a useful connection between the essential set and *quasi-strict* (Nash) equilibria. Quasi-strictness is a refinement of Nash equilibrium that requires that every best response is played with positive probability (Harsanyi 1973). Formally, a Nash equilibrium (σ, t) of a matrix game M is *quasi-strict* if $m(\sigma, j) > m(\sigma, \tau) > m(i, \tau)$ for all $i \in I \setminus supp(\sigma)$ and $j \in J \setminus supp(\tau)$. Since the set of Nash equilibria of a matrix game M is convex, there always exists a Nash equilibrium (σ, τ) with $supp(\sigma) \cup supp(\tau) = ES(M)$. Moreover, it has been shown that all quasi-strict equilibria of a matrix game have the same support (Brandt and Fischer 2008b). Thus, an action is contained in the essential set of a matrix game if and only if it is played with positive probability in some quasi-strict Nash equilibrium. Brandt and Fischer (2008b) have shown that quasi-strict equilibria, and thus the essential set, can be computed in polynomial time.

3.4 Computational Problems

We are interested in the computational complexity of the following decision problems.

- **Possible Equilibrium Action:** Given an incomplete matrix game M and an action a, is there a completion M' of M such that $a \in ES(M')$?
- **Necessary Equilibrium Action:** Given an incomplete matrix game M and an action a, is it the case that $a \in ES(M')$ for all completions M' of M?

One may wonder why these are the right problems to solve. Most generally, the designer could have a utility for each possible outcome (i.e., action profile) of the game. The next proposition shows that hardness of the possible equilibrium action problem immediately implies hardness of the problem of maximizing the designer's utility.

Proposition 1. *Suppose the possible equilibrium action problem is NP-hard. Then, if the designer's payoffs are nonnegative, no positive approximation guarantee for the designer's utility (in the optimistic model where the best equilibrium for the designer is chosen) can be given in polynomial time unless P=NP.*

Proof. Suppose, for the sake of contradiction, that there is a polynomial time algorithm with a positive approximation guarantee for the problem of maximizing the designer's optimistic utility. Then we can use this algorithm for determining whether there is a completion where a strategy receives positive probability in some equilibrium: simply give the designer utility 1 for all outcomes in which that strategy is played, and 0 everywhere else. The designer can get the strategy to be played with positive probability if and only if she can obtain positive utility from this game, and she can obtain positive utility from this game if and only if the approximation algorithm returns a positive utility. □

A similar connection can be given between the necessary equilibrium action problem and the case where designer utilities are nonpositive and a pessimistic model is used (assigning a payoff of -1 to the action in question and 0 otherwise).

In the context of weak tournament games, the essential set (ES) is often interpreted as a (social) choice function identifying desirable alternatives (Dutta and Laslier 1999). In the special case of tournament games, the essential set is referred to as the *bipartisan set* (BP) (Laffond, Laslier, and Le Breton 1993a). The possible and necessary equilibrium action problems defined above thus correspond to possible and necessary *winner* queries for the social choice functions ES (for weak tournament games) and BP (for tournament games). The computational complexity of possible and necessary winners has been studied for many common social choice functions (e.g., Xia and Conitzer 2011; Aziz et al. 2012). To the best of our knowledge, we are the first to provide complexity results for ES and BP.

	c_1	c_2	c_3	c_4	c_5	s_1	s_2	s_3	s_4	s_5
$S_{1,1}$	0	H	$-H$	H	$-H$	y	y	0	0	0
$S_{1,2}$	$\{-1,1\}$	H	$-H$	H	$-H$	x	x	0	0	0
$S_{2,1}$	$-H$	0	H	$-H$	H	0	0	y	y	0
$S_{2,2}$	$-H$	$\{-1,1\}$	H	$-H$	H	0	0	x	x	0
$S_{3,1}$	H	$-H$	0	H	$-H$	0	0	y	0	y
$S_{3,2}$	H	$-H$	$\{-1,1\}$	H	$-H$	0	0	x	0	x
$S_{4,1}$	$-H$	H	$-H$	0	H	0	y	0	y	0
$S_{4,2}$	$-H$	H	$-H$	$\{-1,1\}$	H	0	x	0	x	0
x_1	H	$-H$	H	H	-1	0	0	0	0	0
r_*	$-v$	$-v$	$-v$	$-v$	$-v$	G	G	G	G	G

Figure 1: The incomplete matrix game M used in the proof of Theorem 1 for the SETCOVER instance given by $|U| = 5$, $n = 4, k = 3, S_1 = \{s_1, s_2\}, S_2 = \{s_3, s_4\}, S_3 = \{s_3, s_5\}$, and $S_4 = \{s_2, s_4\}$. L lies in the top left, indicated by double lines.

4 Zero-Sum Games

In this section, we show that computing possible and necessary equilibrium actions is intractable for (not-necessarily-symmetric) matrix games. In the proofs, we will make use of a class of games that we call *alternating games*. Intuitively, an alternating game is a generalized version of Rock-Paper-Scissors that additionally allows "tiebreaking payoffs" which are small payoffs in cases where both players play the same action. A formal definition and proofs of some required properties are given in the full version of this paper.

We first consider the *necessary* equilibrium action problem. Due to the space constraint, and for ease of readability, we only give an informal proof sketch here. Much of the work in the complete proof (to be found in the full version of this paper) is to correctly set values for constants so that the desired equilibrium properties hold.

Theorem 1. *The necessary equilibrium action problem (in matrix games that are not necessarily symmetric) is coNP-complete.*

Proof sketch. For NP-hardness, we give a reduction from SETCOVER. An instance of SETCOVER is given by a collection $\{S_1, \ldots, S_n\}$ of subsets of a universe U, and an integer k; the question is whether we can cover U using only k of the subsets. We may assume that k is odd (it is always possible to add a singleton subset with an element not covered by anything else and increase k by 1). Define an incomplete matrix game M where the row player has $3n - k + 1$ actions, and the column player has $2n - k + |U|$ actions. The row player's actions are given by $\{S_{i,j} : i \in [n], j \in [2]\} \cup \{x_i : i \in [n-k]\} \cup \{r_*\}$.

Let L denote the restriction of the game to the first $2n-k$ columns and $3n - k$ rows. We denote the column player's actions in this part of the game by $c_1, \ldots, c_{2n-k}$. We set $m(S_{i,1}, c_i) = 0$ and $m(S_{i,2}, c_i) = \{-1, 1\}$ for all $i \in [n]$ and $m(x_i, c_{n+i}) = -1$ for all $i \in [n-k]$. We fill in the remaining entries with H and $-H$, where H is a large positive number, so that if we consider only one of each pair of rows $\{S_{i,1}, S_{i,2}\}$, L acts as an alternating game. Setting $m(S_{i,2}, c_i) = -1$ will correspond to choosing S_i for the set cover, and setting $m(S_{i,2}, c_i) = 1$ will correspond to not choosing S_i. Note that, considering only L, the row player will put positive probability on exactly one of $S_{i,1}$ and $S_{i,2}$ (as well as all rows x_i) and, as long as H is sufficiently large, each row that is played with positive probability receives approximately $\frac{1}{N}$ probability. $S_{i,1}$ is played if S_i is chosen for the set cover, $S_{i,2}$ is played otherwise. Also note that the value of L is close to zero, depending on the exact setting of the undetermined entries.

We have additional columns $s_1, \ldots, s_{|U|}$ corresponding to elements of U. For every set S_j containing s_i, column s_i has a positive entry y in row $S_{j,1}$ and a negative entry x in row $S_{j,2}$. If s_i is not covered by any chosen set, then the equilibrium can not be contained in L: If it were, then the column player could best respond by playing s_i, where all entries (on rows played with positive probability by the row player) are either 0 or $x < 0$. However, if s_i is covered by some set, then we can make y large enough (relative to x) that the column player will not play s_i. Thus, if every s_i is covered, the column player plays only columns from L.

Finally, we have a single extra row labeled r_*. This row has a small negative payoff $-v$ for all columns in L, and a very large positive payoff G for all columns not in L. As long as the equilibrium is contained in L (i.e., all elements are covered), it is not a best response for the row player to play r_*. However, if the column player puts positive probability on some s_i (that is, s_i is uncovered), then G is large enough that the row player can best respond by playing r_* with positive probability. □

By modifying the construction in the proof of Theorem 1, we also get a hardness result for the problem of deciding whether an action is a *possible* equilibrium action.

Theorem 2. *The possible equilibrium action problem (in matrix games that are not necessarily symmetric) is NP-complete.*

5 Weak Tournament Games

We now turn to weak tournament games and analyze the computational complexity of possible and necessary *ES* winners.

Theorem 3. *The possible ES winner problem (in weak tournament games) is NP-complete.*

Proof sketch. For NP-hardness, we provide a reduction from SAT. Let $\varphi = C_1 \wedge \ldots \wedge C_m$ be a Boolean formula in conjunctive normal form over a finite set $V = \{v_1, \ldots, v_n\}$ of variables. We define an incomplete weak tournament[2] $W_\varphi = (A, \succ)$ as follows. The set A of vertices is given by $A = \cup_{i=1}^{n} X_i \cup \{c_1, \dots, c_m\} \cup \{d\}$, where $X_i = \{x_i^1, \dots, x_i^6\}$ for all $i \in [n]$. Vertex c_j corresponds to clause C_j and the set X_i corresponds to variable v_i.

Within each set X_i, there is a cycle $x_i^1 \succ x_i^2 \succ x_i^3 \succ x_i^4 \succ x_i^5 \succ x_i^6 \succ x_i^1$ and an unspecified edge between x_i^1 and x_i^4. If variable v_i occurs as a positive literal in clause C_j, we have edges $c_j \succ x_i^3$ and $x_i^5 \succ c_j$. If variable v_i occurs as a negative literal in clause C_j, we have edges $c_j \succ x_i^6$ and $x_i^2 \succ c_j$. Moreover, there is an edge from c_j to d for every $j \in [m]$. For all pairs of vertices for which neither an edge has been defined, nor an unspecified edge declared, we have a tie. See Figure 2 for an example.

We make two observations about W_φ.

Observation 1. *For every completion W of W_φ, we have $d \in ES(W)$ if and only if $ES(W) \cap \{c_1, \dots, c_m\} = \emptyset$.*

Observation 2. *For each i, there is exactly one unspecified edge within (and thus exactly three possible completions of) the subtournament $W_\varphi|_{X_i}$. If the we set a tie between x_i^1 and x_i^4, then all Nash equilibria p of the subtournament $W_\varphi|_{X_i}$ satisfy $p(x_i^1) = p(x_i^3) = p(x_i^5)$ and $p(x_i^2) = p(x_i^4) = p(x_i^6)$. If we set $x_i^1 \succ x_i^4$, then every quasi-strict equilibrium p of $W_\varphi|_{X_i}$ satisfies $p(x_i^2) = p(x_i^4) = p(x_i^6) = 0$, $p(x_i^5) > p(x_i^1) > p(x_i^3) > 0$, and $p(x_i^1) + p(x_i^3) > p(x_i^5)$. By symmetry, setting $x_i^4 \succ x_i^1$ results in quasi-strict equilibria p with $p(x_i^1) = p(x_i^3) = p(x_i^5) = 0$, $p(x_i^4) > p(x_i^6) > p(x_i^2) > 0$, and $p(x_i^2) + p(x_i^6) > p(x_i^4)$.*

We can now show hat φ is satisfiable if and only if there is a completion W of W_φ with $d \in ES(W)$. For the direction from left to right, let α be a satisfying assignment and consider the completion W of W_φ as follows: if v_i is set to "true" under α, add edge $x_i^1 \succ x_i^4$; otherwise, add edge $x_i^4 \succ x_i^1$. It can be shown that $ES(W) = \cup_{i \in [n]} ES(W|_{X_i}) \cup \{d\}$.

For the direction from right to left, let W be a completion of W_φ with $d \in ES(W)$. Define the assignment α by setting variable v_i to "true" if $x_i^1 \succ x_i^4$ and to "false" if $x_i^4 \succ x_i^1$. If there is a tie between x_i^1 and x_i^4, we set the truth value of v_i arbitrarily. Since $d \in ES(W)$, we know by Observation 1 that $c_j \notin ES(W)$ for all $j \in [m]$. It can now be shown that every c_i has an incoming edge from a vertex in $ES(W)$, and that this vertex corresponds to a literal that appears in C_i and that is set to "true" under α. $\square$

We get hardness for the necessary winner problem by slightly modifying the construction used in the proof above.

Theorem 4. *The necessary ES winner problem (in weak tournament games) is coNP-complete.*

It can actually be shown that the problems considered in Theorems 3 and 4 remain intractable even in the case where unspecified payoffs can be chosen from the interval $[-1, 1]$

[2]We utilize the one-to-one correspondence between weak tournament games and directed graphs without cycles of length one or two (so-called *weak tournaments*). For a weak tournament $(A, \succ)$, we use the notation $a \succ b$ to denote a directed edge from a to b.

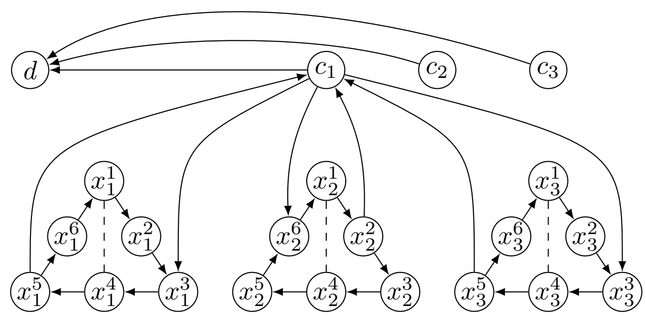

Figure 2: The weak tournament W_φ for formula $\varphi = C_1 \wedge C_2 \wedge C_3$ with $C_1 = x_1 \vee \neg x_2 \vee x_3$. Dashed lines indicate unspecified edges. For improved readability, edges connecting c_2 and c_3 to X have been omitted.

(while still maintaining symmetry). This is interesting insofar as this is our only hardness result for *infinite and/or continuous* payoff sets; such relaxations often make problems computationally easier.

Proposition 2. *The possible equilibrium action problem in weak tournament games remains NP-complete (and the necessary equilibrium action problem coNP-complete) when every payoff set $m(i, j)$ with $m(i, j) = \{-1, 0, 1\}$ is replaced by $m(i, j) = [-1, 1]$.*

6 Tournament Games

The hardness results in Section 5 leave open the possibility that computing ES is tractable in tournament games (where ES is referred to as BP). Indeed, it often turns out that computational problems become easier to solve when restricting attention to tournaments (Baumeister et al. 2013; Brandt and Fischer 2008a; Kenyon-Mathieu and Schudy 2007).[3] The reason is that certain structural properties only hold in tournaments.[4]

Nevertheless, we prove that computing possible and necessary ES winners is hard even in tournament games. The technical difficulty in proving these results lies in the fact that the hardness reduction cannot use "ties" (i.e., non-edges) in (the specified part of) the graph.

Theorem 5. *The possible BP winner problem (in tournament games) is NP-complete.*

Theorem 6. *The necessary BP winner problem (in tournament games) is coNP-complete.*

Observe that these results neither imply the results for weak tournament games in Section 5 (where completions can use "ties") nor the results for general matrix games in Section 4 (where completions can be asymmetric).

[3]There are also cases in the literature where a computational problem remains hard when restricted to tournaments, but the hardness proof is much more complicated (Alon 2006; Charbit, Thomassé, and Yeo 2007; Conitzer 2006).

[4]For example, Laffond, Laslier, and Le Breton (1993b) and Fisher and Ryan (1992) have shown that every tournament game T has a *unique* Nash equilibrium. This Nash equilibrium is quasi-strict and has support $ES(T) = BP(T)$.

7 MIP for Weak Tournament Games

Of course, the fact that a problem is NP-hard does not make it go away; it is still desirable to find algorithms that scale reasonably well (or very well on natural instances). NP-hard problems in game theory often allow such algorithms. In particular, formulating the problem as a *mixed-integer program (MIP)* and calling a general-purpose solver often provides good results. In this section, we formulate the possible ES winner problem in weak tournament games as a MIP.

7.1 Mixed-Integer Programming Formulation

Let $W = (w(i,j))_{i,j \in A}$ be an incomplete weak tournament game. For every entry $w(i,j)$ of W, we define two binary variables x_{ij}^{pos} and x_{ij}^{neg}. Setting $w(i,j)$ to $w_{ij} \in \{-1,0,1\}$ corresponds to setting x_{ij}^{pos} and x_{ij}^{neg} in such a way that $(x_{ij}^{\text{pos}}, x_{ij}^{\text{neg}}) \neq (1,1)$ and $x_{ij}^{\text{pos}} - x_{ij}^{\text{neg}} = w_{ij}$. For each action j, there is a variable p_j corresponding to the probability that the column player assigns to j. Finally, z_{ij} is a variable that, in every feasible solution, equals $w_{ij}p_j$.

To determine whether an action $k \in A$ is a possible ES winner of W, we solve the following MIP. Every feasible solution of this MIP corresponds to a completion of W and a Nash equilibrium of this completion.

maximize p_k

subject to

$$
\begin{array}{ll}
x_{ij}^{\text{neg}} - x_{ji}^{\text{pos}} = 0, \ \forall i,j & x_{ij}^{\text{pos}} = 1, \text{ if } w(i,j) = 1 \\
x_{ij}^{\text{pos}} + x_{ij}^{\text{neg}} \leq 1, \ \forall i,j & x_{ij}^{\text{neg}} = 1, \text{ if } w(i,j) = -1 \\
x_{ij}^{\text{pos}} = x_{ij}^{\text{neg}} = 0, \text{ if } w(i,j) = 0 & x_{ij}^{\text{pos}}, x_{ij}^{\text{neg}} \in \{0,1\}, \ \forall i,j \\
z_{ij} \geq p_j - 2(1 - x_{ij}^{\text{pos}}), \ \forall i,j & \sum_{j \in A} z_{ij} \leq 0, \ \forall i \\
z_{ij} \geq -p_j - 2(1 - x_{ij}^{\text{neg}}), \ \forall i,j & \sum_{j \in A} p_j = 1 \\
z_{ij} \geq -2x_{ij}^{\text{pos}} - 2x_{ij}^{\text{neg}}, \ \forall i,j & p_j \geq 0, \ \forall j
\end{array}
$$

Here, indices i and j range over the set A of actions. Most interesting are the constraints on z_{ij}; we note that exactly one of the three will be binding depending on the values of x_{ij}^{pos} and x_{ij}^{neg}. The net effect of these constraints is to ensure that $z_{ij} \geq w_{ij}p_j$. (Since we also have the constraint $\sum_{j \in A} z_{ij} \leq 0$ and because the value of every completion is zero, $z_{ij} = w_{ij}p_j$ in every feasible solution.) All other constraints containing x_{ij}^{pos} or x_{ij}^{neg} are to impose symmetry and consistency on the entries. The remaining constraints make sure that p is a well-defined probability distribution and that no row yields positive payoff for player 1.

It is possible to adapt this MIP to compute possible and necessary BP winners in tournament games. All that is required is to replace inequality constraints of the form $x_{ij}^{\text{pos}} + x_{ij}^{\text{neg}} \leq 1$ by equalities, thus eliminating the possibility to set $w_{ij} = 0$. Since tournament games have a unique equilibrium, checking whether action k is a possible or necessary BP winner can be done by maximizing and minimizing the objective function p_k, respectively. The reason that this approach does not extend to the computation of necessary winners in *weak* tournament games is that weak tournaments may have multiple equilibria, some of them not quasi-strict. Since our MIP optimizes over the set of all (not necessarily quasi-strict) equilibria, we may encounter cases where the MIP finds a completion with $p_k = 0$, but k is still a necessary winner because it is played with positive probability in every *quasi-strict* equilibrium.

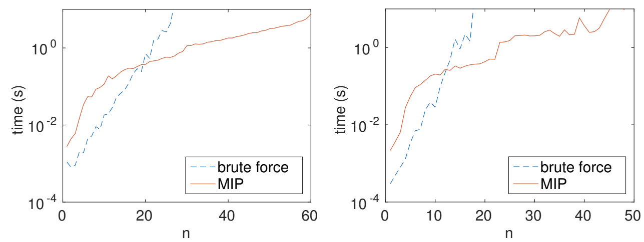

Figure 3: Average runtime (log scale) for $\frac{n}{2}$ unspecified entries (left) and n unspecified entries (right).

7.2 Experimental Results

We tested our MIP for the possible ES winner problem in weak tournament games containing either $\frac{n}{2}$ or n unspecified entries, where $n = |A|$ is the number of actions available to each player. For each n, we examined the average time required to solve 100 random instances[5] of size n, using CPLEX 12.6 to solve the MIP. Results are shown in Figure 3, with algorithms cut off once the average time to find a solution exceeds 10 seconds.

We compared the performance of our MIP with a simple brute force algorithm. The brute force algorithm performs a depth-first search over the space of all completions, terminating when it finds a certificate of a yes instance or after it has exhausted all completions. We observe that for even relatively small values of n, the MIP begins to significantly outperform the brute-force algorithm.

8 Conclusion

Often, a designer has some, but limited, control over the game being played, and wants to exert this control to her advantage. In this paper, we studied how computationally hard it is for the designer to decide whether she can choose payoffs in an incompletely specified game to achieve some goal in equilibrium, and found that this is NP-hard even in quite restricted cases of two-player zero-sum games. Our framework and our results also apply in cases where there is no designer but we are just uncertain about the payoffs, either because further exploration is needed to determine what they are, or because they vary based on conditions (e.g., weather). In such settings one might simply be interested in potential and unavoidable equilibrium outcomes.

Future work may address the following questions. Are there classes of games for which these problems are efficiently solvable? Can we extend the MIP approach to broader classes of games? What results can we obtain for general-sum games? Note that just as hardness for symmetric zero-sum games does not imply hardness for zero-sum

[5]Random instances were generated by randomly choosing each entry from $\{-1,0,1\}$ and imposing symmetry, then randomly choosing the fixed number of entries to be unspecified.

games in general (because in the latter the game does not need to be kept symmetric), in fact hardness for zero-sum games does not imply hardness for general-sum games (because in the latter the game does not need to be kept zero-sum). However, this raises the question of which solution concept should be used—Nash equilibrium, correlated equilibrium, Stackelberg mixed strategies, etc. (All of these coincide in two-player zero-sum games.) All in all, we believe that models where a designer has limited, but not full, control over the game are a particularly natural domain of study for AI researchers and computer scientists in general, due to the problems' inherent computational complexity and potential to address real-world settings.

Acknowledgements

We are thankful for support from NSF under awards IIS-1527434, IIS-0953756, CCF-1101659, and CCF-1337215, ARO under grants W911NF-12-1-0550 and W911NF-11-1-0332, ERC under StG 639945 (ACCORD), a Guggenheim Fellowship, and a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation. This work was done in part while Conitzer was visiting the Simons Institute for the Theory of Computing.

References

Alon, N. 2006. Ranking tournaments. *SIAM Journal on Discrete Mathematics* 20(1):137–142.

An, B.; Shieh, E.; Tambe, M.; Yang, R.; Baldwin, C.; DiRenzo, J.; Maule, B.; and Meyer, G. 2012. PROTECT - A deployed game theoretic system for strategic security allocation for the United States Coast Guard. *AI Magazine* 33(4):96–110.

Anderson, A.; Shoham, Y.; and Altman, A. 2010. Internal implementation. In *Proc. of the 9th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 191–198. IFAAMAS.

Aziz, H.; Brill, M.; Fischer, F.; Harrenstein, P.; Lang, J.; and Seedig, H. G. 2012. Possible and necessary winners of partial tournaments. In *Proc. of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 585–592. IFAAMAS.

Baumeister, D.; Brandt, F.; Fischer, F.; Hoffmann, J.; and Rothe, J. 2013. The complexity of computing minimal unidirectional covering sets. *Theory of Computing Systems* 53(3):467–502.

Bhattacharya, S.; Conitzer, V.; and Munagala, K. 2011. Approximation algorithm for security games with costly resources. In *Proc. of the 7th Workshop on Internet and Network Economics (WINE)*, 13–24.

Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up limit hold'em poker is solved. *Science* 347(6218):145–149.

Brandt, F., and Fischer, F. 2008a. Computing the minimal covering set. *Mathematical Social Sciences* 56(2):254–268.

Brandt, F., and Fischer, F. 2008b. On the hardness and existence of quasi-strict equilibria. In *Proc. of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, 291–302. Springer-Verlag.

Charbit, P.; Thomassé, S.; and Yeo, A. 2007. The minimum feedback arc set problem is NP-hard for tournaments. *Combinatorics, Probability and Computing* 16(1):1–4.

Conitzer, V., and Sandholm, T. 2002. Vote elicitation: Complexity and strategy-proofness. In *Proc. of the 18th National Conference on Artificial Intelligence (AAAI)*, 392–397. AAAI Press.

Conitzer, V. 2006. Computing Slater rankings using similarities among candidates. In *Proc. of the 21st National Conference on Artificial Intelligence (AAAI)*, 613–619. AAAI Press.

Dutta, B., and Laslier, J.-F. 1999. Comparison functions and choice correspondences. *Social Choice and Welfare* 16(4):513–532.

Fisher, D. C., and Ryan, J. 1992. Optimal strategies for a generalized "scissors, paper, and stone" game. *American Mathematical Monthly* 99(10):935–942.

Harsanyi, J. C. 1973. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory* 2:235–250.

Kenyon-Mathieu, C., and Schudy, W. 2007. How to rank with few errors: A PTAS for weighted feedback arc set on tournaments. In *Proc. of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, 95–103. ACM Press.

Konczak, K., and Lang, J. 2005. Voting procedures with incomplete preferences. In *Proc. of the Multidisciplinary Workshop on Advances in Preference Handling*. 124–129.

Laffond, G.; Laslier, J.-F.; and Le Breton, M. 1993a. The bipartisan set of a tournament game. *Games and Economic Behavior* 5(1):182–201.

Laffond, G.; Laslier, J.-F.; and Le Breton, M. 1993b. More on the tournament equilibrium set. *Mathématiques et sciences humaines* 31(123):37–44.

Lang, J.; Pini, M. S.; Rossi, F.; Salvagnin, D.; Venable, K. B.; and Walsh, T. 2012. Winner determination in voting trees with incomplete preferences and weighted votes. *Journal of Autonomous Agents and Multi-Agent Systems* 25(1):130–157.

Monderer, D., and Tennenholtz, M. 2004. K-Implementation. *Journal of Artificial Intelligence Research* 21:37–62.

Pita, J.; Jain, M.; Ordóñez, F.; Portway, C.; Tambe, M.; and Western, C. 2009. Using game theory for Los Angeles airport security. *AI Magazine* 30(1):43–57.

Tsai, J.; Rathi, S.; Kiekintveld, C.; Ordóñez, F.; and Tambe, M. 2009. IRIS - A tool for strategic security allocation in transportation networks. In *Proc. of the 8th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 37–44. IFAAMAS.

Xia, L., and Conitzer, V. 2011. Determining possible and necessary winners under common voting rules given partial orders. *Journal of Artificial Intelligence Research* 41:25–67.

Yin, Z.; Jiang, A. X.; Tambe, M.; Kiekintveld, C.; Leyton-Brown, K.; Sandholm, T.; and Sullivan, J. P. 2012. TRUSTS: Scheduling randomized patrols for fare inspection in transit systems using game theory. *AI Magazine* 33(4):59–72.