

A Computational Analysis of the Tournament Equilibrium Set

Felix Brandt · Felix Fischer · Paul Harrenstein ·
Maximilian Mair

Received: date / Accepted: date

Abstract A recurring theme in the mathematical social sciences is how to select the “most desirable” elements given a binary dominance relation on a set of alternatives. Schwartz’s *tournament equilibrium set* (TEQ) ranks among the most intriguing, but also among the most enigmatic, tournament solutions proposed so far. Due to its unwieldy recursive definition, little is known about TEQ. In particular, its monotonicity remains an open problem to date. Yet, if TEQ were to satisfy monotonicity, it would be a very attractive solution concept refining both the Banks set and Dutta’s minimal covering set. We show that the problem of deciding whether a given alternative is contained in TEQ is NP-hard, and thus does not admit a polynomial-time algorithm unless P equals NP. Furthermore, we propose a heuristic that significantly outperforms the naive algorithm for computing TEQ.

Keywords Tournament Equilibrium Set · Computational Complexity

Mathematics Subject Classification (2000) 91B14 (Social Choice) · 68Q17 (Computational Difficulty of Problems)

JEL classification codes D71 (Social Choice), C63 (Computational Techniques)

1 Introduction

A recurring theme in the mathematical social sciences is how to select the “most desirable” elements given a binary dominance relation on a set of alternatives. Examples are diverse and include selecting socially preferred candidates in social choice settings (e.g., Fishburn, 1977; Laslier, 1997), finding valid arguments in argumentation theory (e.g., Dung, 1995;

F. Brandt · F. Fischer · P. Harrenstein
Institut für Informatik
Ludwig-Maximilians-Universität München, Germany
Tel.: +49-89-2180 9691
Fax: +49-89-2180 9338
E-mail: {brandtf, fischerf, harrenst}@tcs.ifi.lmu.de

Maximilian Mair
Mathematisches Institut
Ludwig-Maximilians-Universität München, Germany
E-mail: mairm@cip.ifi.lmu.de

Dunne, 2007), determining the winners of a sports tournament (e.g., Dutta and Laslier, 1999), making decisions based on multiple criteria (e.g., Bouyssou et al., 2006), choosing the optimal strategy in a symmetric two-player zero-sum game (e.g., Fisher and Ryan, 1995; Laffond et al., 1993b), and singling out acceptable payoff profiles in cooperative game theory (Gillies, 1959; Brandt and Harrenstein, 2009). In social choice theory, where dominance-based solutions are most prevalent, the dominance relation can simply be defined as the pairwise majority relation, i.e., an alternative a is said to dominate another alternative b if the number of individuals preferring a to b exceeds the number of individuals preferring b to a . As is well known from Condorcet's paradox (de Condorcet, 1785), the dominance relation may contain cycles and thus need not have a maximum, even if each of the underlying individual preferences does. As a consequence, the concept of maximality is rendered untenable in most cases, and a variety of so-called *solution concepts* that take over the role of maximality in non-transitive relations have been suggested (see, e.g., Laslier, 1997).

The *tournament equilibrium set* (TEQ) introduced by Schwartz (1990) ranks among the most intriguing, but also among the most enigmatic, solution concepts that has been proposed for tournaments, i.e., asymmetric and complete dominance relations. Due to its unwieldy recursive definition, however, precious little is known about TEQ (Dutta, 1990; Laffond et al., 1993a). In particular, whether TEQ satisfies the important property of monotonicity remains an open question to date. If it does, TEQ constitutes a most attractive tournament solution, refining both the minimal covering set and the Banks set (Laslier, 1997; Laffond et al., 1993a).

The computational effort required to determine a solution is obviously a very important property of any solution concept. If computing a solution is intractable, the applicability of the corresponding solution concept is seriously undermined. This paper relies on the well-established framework of *computational complexity theory* (see, e.g., Papadimitriou, 1994, for an excellent introduction). Complexity theory deals with *complexity classes* of problems that are computationally equivalent in a well-defined way. Typically, problems that can be solved by an algorithm whose running time is polynomial in the size of the problem instance are considered *tractable*, whereas problems that do not admit such an algorithm are deemed *intractable*. The class of decision problems that can be solved in polynomial time is denoted by P, whereas NP (for “nondeterministic polynomial time”) refers to the class of decision problems whose solutions can be *verified* in polynomial time. The famous $P \neq NP$ conjecture states that the hardest problems in NP do not admit polynomial-time algorithms and are thus not contained in P. Although this statement remains unproven, it is widely believed to be true. *Hardness* of a problem for a particular class intuitively means that the problem is no easier than any other problem in that class. Both membership and hardness are established in terms of *reductions* that transform instances of one problem into instances of another problem using computational means appropriate for the complexity class under consideration. In the context of this paper, we will be interested in reductions that can be computed in time polynomial in the size of the problem instances. Finally, a problem is said to be *complete* for a complexity class if it is both contained in and hard for that class. Given the current state of complexity theory, we cannot prove the *actual* intractability of most algorithmic problems, but merely give *evidence* for their intractability. NP-hardness of a problem is commonly regarded as very strong evidence for computational intractability because it relates the problem to a large class of problems for which no efficient, i.e., polynomial-time, algorithm is known, despite enormous efforts to find such algorithms.

In the context of this paper, the definition of any computable tournament solution induces a straightforward algorithm, which exhaustively enumerates all subsets of alternatives and checks which of them comply with the conditions stated in the definition. Not surpris-

ingly, such an algorithm is very inefficient. Yet, proving the intractability of a tournament solution essentially means that *any* algorithm that implements this concept is asymptotically as bad as the straightforward algorithm!

Recent work in computer science has addressed the computational complexity of almost all common tournament solutions (see, e.g., Woeginger, 2003; Alon, 2006; Conitzer, 2006; Brandt et al., 2009). Until recently, the minimal covering set and the tournament equilibrium set have remained two notable exceptions. Laslier writes that “Unfortunately, no algorithm has yet been published for finding the minimal covering set or the tournament equilibrium set of large tournaments. For tournaments of order 10 or more, it is almost impossible to find (in the general case) these sets at hand” (Laslier, 1997, p.8). While the minimal covering set is computable in polynomial time (Brandt and Fischer, 2008), we show that the same is not true for TEQ, unless P equals NP. We first give an arguably simpler alternative to Woeginger’s (2003) NP-hardness proof for membership in the Banks set. Then the construction used in that proof is modified so as to obtain the analogous result for TEQ. In contrast to the Banks set, there is no obvious reason to suppose that the TEQ membership problem is in NP; it may very well be even harder. In the second part of the paper, we propose and evaluate a heuristic for computing TEQ that performs reasonably well on tournaments with up to 150 alternatives. We failed to find a counterexample of TEQ’s conjectured monotonicity by searching a fairly large number of random tournaments.

2 Preliminaries

A *tournament* T is a pair $(A, >)$, where A is a finite set of *alternatives* and $>$ an asymmetric and complete binary relation on A , also referred to as the *dominance relation*. Intuitively, $a > b$ signifies that alternative a beats b in a pairwise comparison.¹ We write \mathcal{T} for the class of all tournaments and have $\mathcal{T}(A)$ denote the set of all tournaments on a fixed set A of alternatives. If T is a tournament on A , then every subset X of A induces a tournament $T|_X = (X, >|_X)$, where $>|_X = \{(x, y) \in X \times X : x > y\}$.

As the dominance relation is not assumed to be transitive in general, there need not be a so-called *Condorcet winner*, i.e., an alternative that dominates all other alternatives. A *tournament solution* S is defined as a function that associates with each tournament T on A a subset $S(T)$ of A . The definition of a tournament solution commonly includes the requirement that $S(T)$ be non-empty if T is defined on a non-empty set of alternatives and that it select the Condorcet winner if there is one (Laslier, 1997, p.37). For X a subset of A , we also write $S(X)$ for the more cumbersome $S(T|_X)$, provided that the tournament T is known from the context. A tournament solution S is said to be *monotonic* if for any two tournaments $T, T' \in \mathcal{T}(A)$ which only differ in that $b > a$ in T and $a > b$ in T' , $a \in S(T)$ implies that also $a \in S(T')$, i.e., reinforcing an alternative cannot cause it to be excluded from the choice set. Monotonicity is a vital property that all reasonable tournament solutions satisfy. In this paper, we will be concerned with two particular tournament solutions, the Banks set and Schwartz’s tournament equilibrium set (TEQ). In order to formally define these concepts, we need some auxiliary notions and notations.

Let R be a binary relation on a set A . We write R^* for the transitive reflexive closure of R . By the *top cycle* $TC_A(R)$ we understand the maximal elements of A according to the asymmetric part of R^* . A subset X of A is said to be *transitive* if R is transitive on X . For

¹ Improving on a previous result by McGarvey, Stearns (1959) has shown that any tournament can be realized via the simple majority rule when the number of voters is at least two greater than the number of alternatives. Thus our results apply to all social choice settings where this is the case.

$X \subseteq Y \subseteq A$, X is called *maximal transitive in Y* if X is transitive and no proper superset of X in Y is. Clearly, since A is finite, every transitive set is contained in a maximal transitive set.

In tournaments, maximal transitive sets are also referred to as Banks trajectories. The *Banks set* $BA(T)$ of a tournament T then collects the maximal elements of the Banks trajectories (Banks, 1985).

Definition 1 (Banks set) Let T be a tournament on A . An alternative $a \in A$ is in the *Banks set* $BA(T)$ of T if a is a maximal element of some maximal transitive set in T .

The *tournament equilibrium set* $TEQ(T)$ of a tournament T on A is defined as the top cycle of a particular subrelation of the dominance relation, referred to as the TEQ relation in the following (Schwartz, 1990). The underlying idea is that alternative a is only “properly” dominated by alternative b , i.e., dominated according to the subrelation, if b is selected from the dominators of a by some tournament solution concept S . To make this idea precise, for $X \subseteq A$, we write $\bar{D}_X(a) = \{b \in X : b > a\}$ for the *dominators* of a in X , omitting the subscript when $X = A$. Thus, for each alternative a one examines the set $\bar{D}(a)$ of its dominators, and solves the subtournament $T|_{\bar{D}(a)}$ by means of the solution S . According to the subrelation, a is then only dominated by the alternatives in $S(\bar{D}(a))$. This of course, still leaves open the question as to the choice of the solution concept S . Now, in the case of TEQ , S is taken to be TEQ itself! This recursion is well-defined because for any $X \subseteq A$ and $a \in X$, the set $\bar{D}_X(a)$ is a *proper* subset of X . Thus, in order to determine the TEQ relation in a subtournament T , one has to calculate the TEQ of subtournaments of T of strictly decreasing order.

Definition 2 (Tournament equilibrium set) Let $T \in \mathcal{T}(A)$. For each subset $X \subseteq A$, define the *tournament equilibrium set* $TEQ(X)$ for X as

$$TEQ(X) = TC_X(\rightarrow_X),$$

where \rightarrow_X is defined as the binary relation on X such that for all $x, y \in X$,

$$x \rightarrow_X y \text{ if and only if } x \in TEQ(\bar{D}_X(y)).$$

Recall that in particular $TEQ(\emptyset) = \emptyset$. The TEQ relation \rightarrow_X is a subset of the dominance relation $>$, and if $\bar{D}_X(x) \neq \emptyset$, then there is some $y \in \bar{D}_X(x)$ with $y \rightarrow_X x$. Furthermore, Definition 2 directly yields a recursive algorithm to compute TEQ. Some reflection reveals that, in the worst case, this *naive algorithm* requires time exponential in $|A|$.

It can easily be established that the Banks set and TEQ both select the Condorcet winner of a tournament if there is one. Moreover, in a cyclic tournament of order three, the Banks set and TEQ both consist of all alternatives. In more complex tournaments, however, the Banks set and TEQ may differ. Consider, for example, the tournament T depicted in Figure 1. We first calculate the TEQ relation \rightarrow . Thus, e.g., for alternative e we find $\bar{D}(e) = \{a, c, d\}$, which constitutes a three-cycle, and so $TEQ(\bar{D}(e)) = \{a, c, d\}$. Accordingly, $a \rightarrow e$, $c \rightarrow e$, as well as $d \rightarrow e$. Doing this for all alternatives, we find $TEQ(T) = \{a, b, c\}$ as the top cycle $TC(\rightarrow)$ of the relation \rightarrow . By contrast, the Banks set consists of the four elements a, b, c and d . For example, $d \in BA(T)$, because $\{d, c, e\}$ is a maximal transitive set with maximal element d . Schwartz (1990) has shown that TEQ is always contained in the Banks set.

Proposition 1 (Schwartz, 1990) Let $T = (A, >)$ be a tournament. Then, $TEQ(T) \subseteq BA(T)$.




Fig. 1 Example due to Schwartz, 1990, where $BA(T) = \{a, b, c, d\}$ and $TEQ(T) = \{a, b, c\}$. The TEQ relation \rightarrow is indicated by thick edges.

Other than that, little is known and much surmised about the theoretical properties of TEQ. For example, Schwartz (1990) conjectured that the top cycle of the TEQ relation always consists of a single connected component, a property of TEQ that is usually referred to as *Index* (Laffond et al., 1993a). Laffond et al. (1993a) and Houy (2009) showed that TEQ satisfying *Index* is equivalent to it having a number of useful properties. In particular, TEQ is monotonic if and only if *Index* holds. Moreover, *Index* implies the inclusion of TEQ in the minimal covering set, another appealing tournament solution. Thus, if TEQ satisfies *Index* it might be considered a very strong solution concept. Otherwise, TEQ lacks the vital property of monotonicity and as such it would be severely flawed as a tournament solution.

3 An Alternative NP-Hardness Proof for Membership in the Banks Set

We begin our investigation of the computational complexity of the TEQ membership problem by giving an alternative proof for NP-hardness of the analogous problem for the Banks set. The latter was first demonstrated by Woeginger (2003) using a reduction from graph three-colorability. Our proof works by a reduction from 3SAT, the NP-complete satisfiability problem for Boolean formulas in conjunctive normal form with exactly three literals per clause (see, e.g., Papadimitriou, 1994). It is arguably simpler than Woeginger's, and a much similar construction will be used in the next section to prove NP-hardness of membership in TEQ. The tournaments used in both reductions will be taken from a special class \mathcal{T}^* .

Definition 3 (The class \mathcal{T}^*) A tournament $(A, >)$ is in the class \mathcal{T}^* if it satisfies the following properties. There is some odd integer $n \geq 1$, the number of *layers* in the tournament, such that $A = C \cup U_1 \cup \dots \cup U_n$, where C, U_1, \dots, U_n are pairwise disjoint and $C = \{c_0, \dots, c_n\}$. Each U_i is a singleton if i is even, and $U_i = \{u_i^1, u_i^2, u_i^3\}$ if i is odd. The dominance relation $>$ satisfies the following properties for all $c_i \in C_i, c_j \in C_j, u_i \in U_i, u_j \in U_j$ ($0 \leq i, j \leq n$):

- (i) $c_i > c_j$, if $i > j$,
- (ii) $u_i > c_j$, if $i = j$,
- (iii) $c_j > u_i$, if $i \neq j$,
- (iv) $u_i > u_j$, if $i < j$ and at least one of i and j is even,
- (v) $u_i^k > u_i^l$, if i is odd and $k \equiv l - 1 \pmod{3}$.

We also refer to c_0 by d , for “decision node” and to $\bigcup_{1 \leq i \leq n} U_i$ by U . For $i = 2k$, we have as a notational convention $U_i = Y_k = \{y_k\}$ and set $Y = \bigcup_{1 \leq 2k \leq n} Y_k$.

Observe that this definition fixes the dominance relation between any two alternatives except for some pairs of alternatives that are both in U .

As a next step in the argument, we associate with each instance of 3SAT a tournament in the class \mathcal{T}^* . An instance of 3SAT is given by a formula φ in 3-conjunctive




Fig. 2 Tournament T_φ^{BA} for the 3CNF formula $\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r)$. Omitted edges are assumed to point downwards.

normal form (3CNF), i.e., $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \dots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$, where each $x \in \{x_i^1, x_i^2, x_i^3 : 1 \leq i \leq m\}$ is a literal. For each clause $x_i^1 \vee x_i^2 \vee x_i^3$ we assume x_i^1, x_i^2 and x_i^3 to be distinct literals. We moreover assume the literals to be indexed and by X_i we denote the set $\{x_i^1, x_i^2, x_i^3\}$. For literals x we have $\bar{x} = \neg p$ if $x = p$, and $\bar{x} = p$ if $x = \neg p$, where p is some propositional variable. We may also assume that if x and y are literals in the same clause, then $x \neq \bar{y}$. We say a 3CNF $\varphi = (x_1^1 \vee x_1^2 \vee x_1^3) \wedge \dots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$ is *satisfiable* if there is a tuple (x_1, \dots, x_m) in $\times_{1 \leq i \leq m} X_i$ such that $v' = \bar{v}$ for no $v, v' \in \{x_1, \dots, x_m\}$. Let $V = \{x_1, \dots, x_m\}$. Next we define for each 3SAT formula φ the tournament T_φ^{BA} .

Definition 4 (Banks construction) Let φ be a 3CNF $(x_1^1 \vee x_1^2 \vee x_1^3) \wedge \dots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$. Define $T_\varphi^{BA} = (C \cup U, >)$ as the tournament in the class \mathcal{T}^* with $2m - 1$ layers such that for all $1 \leq j < 2m$,

$$U_j = \begin{cases} X_i & \text{if } j = 2i - 1, \\ \{y_i\} & \text{if } j = 2i \end{cases}$$

and such that for all $x \in X_i$ and $x' \in X_j$ ($1 \leq i, j \leq m$),

$$x > x' \quad \text{if both } j < i \text{ and } x' = \bar{x} \text{ or both } i < j \text{ and } x' \neq \bar{x}.$$

Observe that in conjunction with the other requirements on the dominance relation of a tournament in \mathcal{T}^* , this completely fixes the dominance relation $>$ of T_φ^{BA} .

An example of a tournament T_φ^{BA} for a 3CNF φ is shown in Figure 2. We are now in a position to present our proof that the Banks membership problem is NP-complete.

Theorem 1 *The problem of deciding whether a particular alternative is in the Banks set of a tournament is NP-complete.*

Proof Membership in NP is obvious. For any fixed alternative a , we can simply guess a transitive subset of alternatives V with a as maximal element and verify that V is also maximal with respect to set inclusion.

For NP-hardness, we show that T_φ^{BA} contains a maximal transitive set with maximal element d if and only if φ is satisfiable. First observe that V is a maximal transitive subset with maximal element d in T_φ^{BA} only if both

- (i) for all $1 \leq i < 2m$ there is a $u \in U_i$ such that $u \in V$, and
- (ii) there are no $1 \leq i < j < 2m$, $u \in U_i$, $u' \in U_j$ with $u, u' \in V$ such that $u' > u$.

Regarding (i), if there is an $1 \leq i < 2m$ such that no element of U_i is contained in V , we can always add c_i to V in order to obtain a larger transitive set. If (ii) were not to hold, both i and j have to be odd for u_j to dominate u_i . However, in light of (i), there has to be a k with $i < k < j$ and $u'' \in U_k$ such that $u'' \in V$. It follows that V is not transitive because u, u'' , and u' form a cycle. Recall we may assume that $x \neq \bar{y}$ if x and y are literals in the same clause. Thus, if there is maximal transitive set V with maximal element d complying with both (i) and (ii), a satisfying assignment of φ can be obtained by letting all literals contained in $X \cap V$ be true.

For the opposite direction, assume that φ is satisfiable. Then there is a tuple (x_1, \dots, x_m) in $\times_{1 \leq i \leq m} X_i$ such that $x' = \bar{x}$ for no $x, x' \in \{x_1, \dots, x_m\}$. Obviously $V = \{x_1, \dots, x_m\} \cup \{y_1, \dots, y_{m-1}\} \cup \{d\}$ contains no cycles and thus is transitive with maximal element d . In order to obtain a larger transitive set with a different maximal element, we need to add c_i for some $1 \leq i \leq m$ to V . However, $V \cup \{c_i\}$ always contains a cycle consisting of c_i , d , and u for some $u \in U_i$, contradicting the transitivity of $V \cup \{c_i\}$. We have thus shown that d is the maximal element of some maximal transitive set in T_φ^{BA} containing V as a subset. \square

4 NP-hardness of Membership in TEQ

In this section we prove that the problem of deciding whether a particular alternative is in the TEQ of a tournament is NP-hard. To this end, we refine the construction that was used in the previous section to prove NP-completeness of membership in the Banks set.

Definition 5 (TEQ construction) Let φ be a 3CNF $(x_1^1 \vee x_1^2 \vee x_1^3) \wedge \dots \wedge (x_m^1 \vee x_m^2 \vee x_m^3)$. Further for each $1 \leq i < m$, let there be a set $Z_i = \{z_i^1, z_i^2, z_i^3\}$. Define T_φ^{TEQ} as the tournament $(A, >)$ in \mathcal{T}^* with $4m - 3$ layers such that $A = C \cup U_1 \cup \dots \cup U_{4m-3}$ and for all $1 \leq i \leq m$,

$$U_j = \begin{cases} X_i & \text{if } j = 4i - 3, \\ Z_i & \text{if } j = 4i - 1, \\ \{y_i\} & \text{otherwise.} \end{cases}$$

As in the Banks construction, we let for all $x \in X_i$ and $x' \in X_j$ ($1 \leq i, j \leq m$)

$$x > x' \quad \text{if both } j < i \text{ and } x' = \bar{x} \text{ or both } i < j \text{ and } x' \neq \bar{x}.$$

Finally, for all $1 \leq i, j \leq m$, $x_i^k \in X_i$ and $z_j^l \in Z_j$,

$$x_i^k > z_j^l \text{ if and only if } i < j \text{ or both } i = j \text{ and } k = l.$$

An example for such a tournament is shown in Figure 3.

We now proceed to show that a 3SAT formula φ is satisfiable if and only if the decision node d is in the tournament equilibrium set of T_φ^{TEQ} . We make use of the following lemma. Recall that \rightarrow_B^* denotes the transitive reflexive closure of \rightarrow_B . Moreover, for $B \subseteq A$ with $B \cap C \neq \emptyset$, let c_B be the alternative in C with the highest index among those included in B .




Fig. 3 Tournament T_φ^{TEQ} for the 3CNF formula $\varphi = (\neg p \vee s \vee q) \wedge (p \vee s \vee r) \wedge (p \vee q \vee \neg r)$.

Lemma 1 Let $T = (A, >)$ be a tournament in \mathcal{T}^* and let $B \subseteq A$ such that $d \in B$. Then, $c_B \rightarrow_B^* b$ for all $b \in B$.

Proof The proof is by induction on the size of B . The basis, i.e., if $B = \{d\}$, is trivial. Now assume that $\{d\}$ is a proper subset of B and consider an arbitrary $b \in B$. First assume $b \in B \cap C$. If $b \neq c_B$, c_B is the Condorcet winner in $\overline{D}_B(b)$ by construction. Hence, either $c_B = b$ or $c_B \rightarrow_B b$. In either case $c_B \rightarrow_B^* b$. Now assume $b \in B \cap U$ and consider $\overline{D}_B(b)$. As $\overline{D}_B(b)$ is a proper subset of B including d the induction hypothesis is applicable. Hence, $c_{\overline{D}_B(b)} \rightarrow^* x$ for all $x \in \overline{D}_B(b)$. Accordingly, $c_{\overline{D}_B(b)} \in TEQ(\overline{D}_B(b))$ and thus $c_{\overline{D}_B(b)} \rightarrow_B b$. As also $c_B \rightarrow_B^* c_{\overline{D}_B(b)}$, it follows that $c_B \rightarrow_B^* b$. \square

We are now ready to state the main theorem of this paper.

Theorem 2 Deciding whether a particular alternative is in the tournament equilibrium set of a tournament is NP-hard.

Proof By reduction from 3SAT. Consider an arbitrary 3CNF φ and construct the tournament $T_\varphi^{TEQ} = (C \cup U, >)$. This can be done in polynomial time. We show that

$$\varphi \text{ is satisfiable if and only if } d \in TEQ(T_\varphi^{TEQ}).$$

For the direction from right to left, observe that, by an argument analogous to the proof of Theorem 1, it can be shown that φ is satisfiable if and only if $d \in BA(T_\varphi^{TEQ})$. So assuming

that φ is not satisfiable yields $d \notin BA(T_\varphi^{TEQ})$. By the inclusion of TEQ in the Banks set (see Proposition 1) it follows that $d \notin TEQ(T_\varphi^{TEQ})$.

For the opposite direction, assume that φ is satisfiable. Then there is a tuple (x_1, \dots, x_m) in $\times_{1 \leq i \leq m} X_i$ such that $x' = \bar{x}$ for no $x, x' \in \{x_1, \dots, x_m\}$. Let $W = \{x_1, \dots, x_m\}$ and $\{u_1, \dots, u_n\} = W \cup \{y_1, \dots, y_{m-1}\} \cup \{z_i^j \in Z: x_i^j \in W\}$, where $u_i \in U_i$ for each $1 \leq i \leq n$. Obviously, $\{u_1, \dots, u_n\}$ contains no cycles and as such is transitive. We now define a sequence B_1, \dots, B_{n+1} of subsets of A such that for each i with $1 \leq i \leq n+1$

$$B_i = \begin{cases} A & \text{if } i = n+1, \\ \overline{D}(u_i) \cap B_{i+1} & \text{otherwise.} \end{cases}$$

Defined thus we have $B_1 \subseteq \dots \subseteq B_{n+1}$. Moreover, by construction, $d \in B_i$ for each i with $1 \leq i \leq n+1$. To simplify notation, we write \rightarrow_i and $\overline{D}_i(x)$ for \rightarrow_{B_i} and $\overline{D}_{B_i}(x)$, respectively. Observe that for all i with $1 \leq i \leq n$,

$$B_i = \overline{D}_{i+1}(u_i) = \overline{D}(u_i) \cap \dots \cap \overline{D}(u_n).$$

It now suffices to prove that

$$d \in TEQ(B_k) \text{ for all } 1 \leq k \leq n+1. \quad (*)$$

We first make the following observations. For all $1 \leq i \leq n+1$ and $1 \leq j \leq n$,

- (i) $u_j \in B_i$ if and only if $j < i$,
- (ii) $c_j \in B_i$ if and only if $j < i$,
- (iii) $c_i \rightarrow_{i+1} c_j$ if $j < i \leq n$,
- (iv) $u_i \rightarrow_{i+1} c_i$, if $i \leq n$.

For (i), observe that $u_j \in A$ and, by transitivity of the set $\{u_1, \dots, u_n\}$, $u_j \in \overline{D}(u_i)$, \dots , $u_j \in \overline{D}(u_n)$ if $i \leq n$. Hence, $u_j \in B_i$. If on the other hand $j \geq i$, then $u_j \notin \overline{D}(u_j)$ and thus $u_j \notin B_i$. For (ii), observe that $c_j \in \overline{D}(u_i)$ for all $i \neq j$ and thus $c_j \in B_i$ if $j < i$. However, $c_j \notin \overline{D}(u_j)$ and hence $c_j \notin B_i$ if $j \geq i$. For (iii), merely observe that c_i is the Condorcet winner in $\overline{D}_{i+1}(c_j)$, if $j < i \leq n$. To appreciate (iv), observe that by construction $\overline{D}_{i+1}(c_i)$ has to be either a singleton $\{u_i\}$ for some $u_i \in U_i$, or U_i itself. The former holds if $U_i \subseteq Y$, or if $U_i \subseteq X$ and $i \neq n$, the latter if $U_i = U_n$ or if $U_i \subseteq Z$. In either case, $TEQ(\overline{D}_{i+1}(c_i)) = \overline{D}_{i+1}(c_i)$ and $u_i \rightarrow_{i+1} c_i$ holds. To see that $\overline{D}_{i+1}(c_i) = \{u_i\}$ if $U_i \subseteq X$ with $i \neq n$, let $U_i = \{u_i, u'_i, u''_i\}$. By construction, $\overline{D}(c_i) = U_i \cup \{c_j: j > i\}$ and $U_{i+2} \subseteq Z$. Moreover, by transitivity of $\{u_1, \dots, u_n\}$, also $u'_i, u''_i \notin \overline{D}(u_{i+2})$. Accordingly, $u'_i, u''_i \notin B_{i+1}$ whereas, with (i), we have $u_i \in B_{i+1}$. In virtue of (ii), we may conclude that $\overline{D}_{i+1}(c_i) = \overline{D}(c_i) \cap B_{i+1} = \{u_i\}$.

We are now in a position to prove (*) by induction on k . For $k = 1$, observe that, by construction and observation (ii), d is a Condorcet winner in B_1 and, thus, $d \in TEQ(B_1)$. For the induction step, let $k = i+1$. By observation (i) we know that $u_i \in B_{i+1}$ and, in virtue of the induction hypothesis, that $d \in TEQ(B_i)$. Because $B_i = \overline{D}_{i+1}(u_i)$, we have $d \rightarrow_{i+1} u_i$. Moreover, by observations (iii) and (iv),

$$c_i \rightarrow_{i+1} d \rightarrow_{i+1} u_i \rightarrow_{i+1} c_i,$$

i.e., c_i, d and u_i constitute a \rightarrow_{i+1} -cycle. In virtue of observation (ii), c_i is the alternative in $C \cap B_{i+1}$ with the highest index. By Lemma 1 it then follows that $c_i \rightarrow_{i+1}^* b$ for all $b \in B_{i+1}$. We may conclude that $\{c_i, d, u_i\} \subseteq TC_{B_{i+1}}(\rightarrow_{i+1})$. Hence, $d \in TEQ(B_{i+1})$. \square

The preceding proof establishes the intractability of *any* solution concept that is sandwiched between the Banks set and TEQ, i.e., any concept that always selects a subset of the Banks set and a superset of TEQ.

Algorithm 1 Tournament Equilibrium Set

```

procedure TEQ( $X$ )
 $R \leftarrow \emptyset$ 
 $B \leftarrow C \leftarrow \arg \min_{a \in X} |\overline{D}(a)|$ 
loop
   $R \leftarrow R \cup \{(b, a) : a \in C \wedge b \in \text{TEQ}(\overline{D}(a))\}$ 
   $D \leftarrow \bigcup_{a \in C} \text{TEQ}(\overline{D}(a))$ 
  if  $D \subseteq B$  then return  $TC_B(R)$  end if
   $C \leftarrow D$ 
   $B \leftarrow B \cup C$ 
end loop

```

5 A Heuristic for Computing TEQ

Computational intractability of the TEQ membership problem implies that TEQ cannot be computed efficiently either. Nevertheless, the running time of the naive algorithm, which straightforwardly implements the recursive definition of TEQ, can be greatly reduced when assuming that TEQ satisfies *Index*. This assumption can fairly be made because otherwise TEQ would be severely compromised as a solution concept and the issue of computing it moot.

Algorithm 1 computes TEQ by starting with the set B of all alternatives that have dominator sets of minimal size (i.e., the so-called Copeland winners). These alternatives are good candidates to be included in TEQ, and the small size of their dominator sets speeds up the computation of their TEQ-dominators. Then, all alternatives that TEQ-dominate any alternative in B are iteratively added to B . When no more such alternatives can be found, the algorithm returns the top cycle of \rightarrow_B . The *worst-case* running time of this algorithm is of course still exponential, but experimental results suggest that it significantly outperforms the naive algorithm on tournaments that are generated by orienting each edge uniformly at random. These results are shown in Table 1.

We implemented different versions of both algorithms, based on different subroutines to determine the top cycle. Among these, the algorithm by Tarjan (1972) consistently showed the best performance. The recursive computation of TEQ involves the computation of the TEQ of many tournaments of low order. A significant speedup can therefore be achieved by computing the TEQ only once for small tournaments and storing the result in a lookup table. In practice, the effectiveness of this approach is limited by the amount of memory available to store the results. In our experiments, a good tradeoff between computation time and memory requirements was obtained by storing the TEQ for tournaments of order nine or less.

While choosing tournaments uniformly at random might be useful for benchmarking algorithms, it raises a number of conceptual problems. First, in voting and most other applications uniform random tournaments do not represent a reasonably realistic model of social preferences. Secondly, these tournaments are “almost” regular and most tournament solutions almost always select all alternatives in regular tournaments. One model of random tournaments that have more structure can be obtained by defining an arbitrary linear order on the alternatives a_1, \dots, a_n and letting $a_i > a_j$ for $i < j$ with some fixed probability $p > 0.5$. The case where $p = 1$ yields a “completely structured” transitive tournament. The more structure a tournament possesses, the more Algorithm 1 outperforms the naive algorithm, due to the increasing number of large dominator sets that have to be analyzed by the latter at every level of the recursion. In large structured tournaments, the performance gap becomes

A	naive		Algorithm 1	
	w/o lookup	w/ lookup	w/o lookup	w/ lookup
Uniform random tournaments ($p = 0.5$)				
50	0.65 s	0.28 s	0.19 s	0.18 s
100	49.60 s	9.67 s	9.28 s	5.84 s
150	1 088 s	178 s	196 s	105 s
200	10 525 s	1 731 s	1 910 s	1 034 s
Structured random tournaments ($p = 0.8$)				
50	19.24 s	2.02 s	0.02 s	0.02 s
100	15 981 s	1 504 s	3.03 s	2.49 s
150	> 24 h	> 24 h	720 s	522 s
200	> 24 h	> 24 h	44 454 s	31 197 s

Table 1 Experimental evaluation of algorithms for computing TEQ. Both the naive algorithm and Algorithm 1 were implemented in the Java programming language. The table lists the average running time for ten instances on a 3GHz Core2Duo machine. Running times are also given for versions of the algorithms that store the TEQ of tournaments of order nine or less in a lookup table (which is initially empty). For tournaments of order 150 or more, the naive algorithm did not terminate within 24 hours.

rather impressive (see Table 1). For example, to compute the TEQ of a structured random tournament of order 100, the naive algorithm requires more than four hours, or about 25 minutes with lookup, whereas this takes Algorithm 1 only about three seconds.²

We have further used the naive algorithm to try to disprove *Index* (and thus TEQ’s monotonicity), but failed to find a counterexample by an exhaustive search in all tournaments of order 12 or lower. Over a period of several months we also investigated about 8 billion random tournaments with up to 50 alternatives, again to no avail. This is interesting insofar as we successfully used the same methods to find counterexamples for TEQ’s external stability, a related conjecture by Schwartz that was recently disproved by Houy (2009). We found plenty of counterexamples for external stability by random search and showed that Houy’s counterexample of 11 alternatives is minimal. The total number of counterexamples of order 11 is 26.

Acknowledgements This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/3-2. We are grateful to Brendan McKay for providing a program that generates non-isomorphic tournaments.

References

- N. Alon. Ranking tournaments. *SIAM Journal on Discrete Mathematics*, 20(1):137–142, 2006.
- J. S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 3:295–306, 1985.
- D. Bouyssou, T. Marchant, M. Pirlot, A. Tsoukiàs, and P. Vincke. *Evaluation and Decision Models: Stepping Stones for the Analyst*. Springer-Verlag, 2006.
- F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.
- F. Brandt and P. Harrenstein. Characterization of dominance relations in finite coalitional games. *Theory and Decision*, 2009. To appear.

² We also had some limited success with algorithms that make use of the easy-to-prove fact that $TEQ((A, >)) = TEQ(TC_A(>))$. Assuming *Index*, a similar preprocessing step that first computes the minimal covering set of the tournament at hand is possible.

- F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. *Mathematical Logic Quarterly*, 55(4):444–459, 2009.
- V. Conitzer. Computing Slater rankings using similarities among candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 613–619. AAAI Press, 2006.
- Marquis de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.
- P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- P. E. Dunne. Computational properties of argumentation systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171(10-15):701–729, 2007.
- B. Dutta. On the tournament equilibrium set. *Social Choice and Welfare*, 7(4):381–383, 1990.
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
- P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
- D. C. Fisher and J. Ryan. Tournament games and positive tournaments. *Journal of Graph Theory*, 19(2):217–236, 1995.
- D. B. Gillies. Solutions to general non-zero-sum games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games IV*, volume 40 of *Annals of Mathematics Studies*, pages 47–85. Princeton University Press, 1959.
- N. Houy. Still more on the tournament equilibrium set. *Social Choice and Welfare*, 32: 93–99, 2009.
- G. Laffond, J.-F. Laslier, and M. Le Breton. More on the tournament equilibrium set. *Mathématiques et sciences humaines*, 31(123):37–44, 1993a.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5:182–201, 1993b.
- J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.
- C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- T. Schwartz. Cyclic tournaments and cooperative majority voting: A solution. *Social Choice and Welfare*, 7:19–29, 1990.
- R. Stearns. The voting problem. *American Mathematical Monthly*, 66(9):761–763, 1959.
- R. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972.
- G. J. Woeginger. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20:523–528, 2003.